arXiv:1604.05703v1 [math.PR] 19 Apr 2016

A large deviations analysis of certain qualitative properties of parallel tempering and infinite swapping algorithms

J. Doll*, P. Dupuis† and P. Nyquist‡

April 20, 2016

Abstract

Parallel tempering, or replica exchange, is a popular method for simulating complex systems. The idea is to run parallel simulations at different temperatures, and at a given swap rate exchange configurations between the parallel simulations. From the perspective of large deviations it is optimal to let the swap rate tend to infinity and it is possible to construct a corresponding simulation scheme, known as infinite swapping. In this paper we propose a novel use of large deviations for empirical measures for a more detailed analysis of the infinite swapping limit in the setting of continuous time jump Markov processes. Using the large deviations rate function and associated stochastic control problems we consider a diagnostic based on temperature assignments, which can be easily computed during a simulation. We show that the convergence of this diagnostic to its a priori known limit is a necessary condition for the convergence of infinite swapping. The rate function is also used to investigate the impact of asymmetries in the underlying potential landscape, and where in the state space poor sampling is most likely to occur.

*Department of Chemistry, Brown University, Providence, RI 02912. Research supported in part by the National Science Foundation (DMS-1317199), and the Defense Advanced Research Projects Agency (W911NF-15-2-0122).

†Division of Applied Mathematics, Brown University, Providence, RI 02912. Research supported in part by the Department of Energy (de-sc0010539), the National Science Foundation (DMS-1317199), and the Defense Advanced Research Projects Agency (W911NF-15-2-0122).

‡Division of Applied Mathematics, Brown University, Providence, RI 02912. Research supported in part by National Science Foundation (DMS-1317199).

1 Introduction

In prior work [6] we used the large deviation rate function to study rate of convergence questions for parallel tempering (PT, also known as replica exchange) type computational methods. The analysis suggested the construction of a related method obtained in the limit as the attempted swap rate tends to infinity, which was labeled the infinite swapping (INS) algorithm, along with more generally implementable partial infinite swapping (PINS) algorithms. In the present paper we apply the rate function again along with related constructions from stochastic optimal control to analyze certain qualitative properties of INS and, to a lesser extent, of PT algorithms.

There are two main results. The first is the theoretical analysis of a diagnostic for good sampling that was introduced in [2, 1]. The empirical measure of particle/temperature associations (i.e., the fraction of time that a given particle is assigned to a given dynamic) is a quantity that is easy to record during a simulation. It was shown in the references that under mild assumptions, for PT, PINS and INS the empirical measure converges to the uniform distribution in the large time limit. We show in this paper that if the empirical measure is not in a small neighborhood of the uniform distribution, then with overwhelming probability the numerical approximation provided by INS is not close to the target product measure. Hence the particle/temperature empirical measure provides a convenient diagnostic for good sampling. Such a diagnostic can be very useful when applying Monte Carlo to problems involving rare events, where the approximation will in general not converge in a gradual and predictable way, and other more standard diagnostics (e.g., empirical variance) may suggest convergence when it has not occurred. Of course temperature is only one of several parameters that can (and have) been used to index the ensemble of systems used in parallel replica methods, and we expect that the results developed here for temperature could be generalized to these other parameters as well [7].

The second use of large deviation theory is to study other qualitative properties of INS. It is known that the large deviation rate function can tell one not just the decay rates associated with rare events, but also the most likely way that the event will occur. This property has found many uses in the application of large deviations to problems involving small random perturbations of deterministic systems, where it is used to indicate those pathways a process is most likely to follow on the way to a rare outcome. However, our use here is fundamentally different, in that we will use the large deviation rate function for empirical measures to understand and explain various properties of sampling schemes. We consider two such appli-

cations. The first is to characterize the impact of symmetries in the energy landscape on the convergence properties of INS (and also to some degree PT). Indeed, we find that symmetry properties of the potential play a large role in the performance of PT and INS schemes, and uncover some counter-intuitive behavior. For example, we show that reducing energy barriers can actually slow convergence when the reduction leads to a type of symmetry breaking. The second application is to identify those parts of the state space where "errors" in the sampling have the greatest impact on the overall performance of the algorithms. Such information could, in principle, be used to design a sampling scheme which focuses computational effort on accurate approximation of distributional properties in these regions, thereby leading to better performance. Both of these applications exploit the fact that we solve variational problems associated with the rate function by converting them to ergodic (average cost per unit time) stochastic optimal control problems, which are then solved numerically. To the best of our knowledge this is the first instance where large deviations results have been used in this way, allowing us to obtain explicit information on the convergence properties of simulation schemes.

The organization of the paper is as follows. Section 2 introduces the INS algorithm for a general class of jump Markov processes. The use of this particular type of process is motivated in part by the fact that variational problems involving the rate function can be solved numerically. The particle/temperature empirical measure is also introduced. Section 3 describes the joint large deviation properties of the INS numerical approximation and the particle/temperature empirical measure. The extraction of information from the rate function will require the solution to certain variational problems, for which a conversion into equivalent ergodic stochastic control problems is more convenient. These controls problems and their properties are given in Section 4. Although the material of this section is known (at least in some form) in the literature on stochastic control, we could not find it in the form we need, nor is it likely to be familiar to those using PT and related methods. We then apply these ideas to the analysis of the diagnostic in Section 5, and the analysis of other qualitative properties in Section 6.

2 Infinite swapping for jump Markov processes

In this section we describe the infinite swapping algorithms that are appropriate for discrete spaces. Specifically, we consider the setting of a single temperature continuous-time Markov jump processes on a finite state space

$\mathcal{S}$ of size $|\mathcal{S}| = N \in \mathbb{N}$. For a description of infinite swapping corresponding to diffusion models or jump Markov processes on an uncountable state space see [6]. The results we present will carry over to these settings but with a more involved analysis. Also, as noted previously the discrete formulation makes the numerical solution of examples much simpler.

The infinite swapping algorithms take values in the product space $\mathcal{S}^K$, where $K \in \mathbb{N}$ is the number of temperatures under consideration. For $K \in \mathbb{N}$, x, y, z and so on are used to denote generic elements of $\mathcal{S}$, and the boldface versions such as $\mathbf{x} = (x_1, x_2, \ldots, x_K)$ denote elements of the product space. For notational brevity we often consider what corresponds to a two temperature model and typically only comment on the extension to multiple temperatures. It will become clear that, from a mathematical viewpoint, the inclusion of more temperatures is much a matter of bookkeeping. However, the use of many temperatures eventually introduces a practical challenge when implementing infinite swapping algorithms. This is addressed in [6] where the partial infinite swapping (PINS) algorithm is introduced.

In the setting of Markov jump processes with two temperatures τ_1 and τ_2, we describe the dynamics of the process by two rate matrices, Γ^1 and Γ^2, each of size $N \times N$. One should think of Γ^1 as the dynamics that go with τ_1 and similarly for Γ^2 and τ_2. For two states $x \neq y$, $\Gamma^i_{x,y} \geq 0$ represents the rate at which the process jumps from x to y.

As an example of rate matrices Γ^1 and Γ^2 that correspond to temperatures τ_1 and τ_2, we consider a particular form of so-called Glauber dynamics [15]. To be precise, let μ be a Gibbs measure defined in terms of a potential $V : \mathcal{S} \to \mathbb{R}$ and temperature τ:

$$\mu(x) = e^{-V(x)/\tau}/Z(\tau),\ x \in \mathcal{S},$$

where $Z(\tau)$ is the normalizing constant associated with the temperature τ.

To define a type of Glauber dynamics with μ as invariant measure, let A be an $N \times N$ matrix with entries $A_{x,y} \in \{0, 1\}$ for all $x, y \in \mathcal{S}$, with $A_{x,x} = 0$. The matrix A defines the communicating class(es) of $\mathcal{S}$. Next, we define a rate matrix Γ in terms of the potential V and the temperature τ by

$$\begin{aligned}
\Gamma_{x,y} &\doteq \exp\left\{-\frac{1}{\tau}\left(V(y) - V(x)\right)^+\right\} A_{x,y},\ y \neq x,\\
\Gamma_{x,x} &\doteq -\sum_{y \neq x} \Gamma_{x,y}.
\end{aligned}$$

Then $q(x) = -\Gamma_{x,x}$ represents the total rate out of state x for a chain with dynamics according to Γ. Note that this is but one particular form of Glauber

dynamics, and is sometimes referred to as Metropolis dynamics; see [15] for other forms and details. For any such Glauber dynamics, μ is an invariant distribution for the associated continuous time Markov chain. Under ergodicity, the empirical measure for such a chain will converge to μ. We assume Γ^1 and Γ^2 are ergodic henceforth.

2.1 Infinite swapping limit

For simplicity we continue to keep the discussion to two temperatures τ_1 and τ_2. In accordance with the previous section, let μ_1 and μ_2 denote the invariant distributions associated with the given dynamics and the two rate matrices be denoted Γ^1 and Γ^2, respectively. Let μ be the product distribution

$$\mu = \mu_1 \times \mu_2.$$

The infinite swapping limit for this setting is described in [6] and we review only the basics. Let $\mathbf{X}^0 = \{(X_1^0(t), X_2^0(t)) : t \geq 0\}$ denote the Markov process with independent components, each having dynamics according to the rate matrices Γ^i, $i = 1, 2$. The embedded Markov chain $\bar{\mathbf{X}}^0$ has probability transition kernel

$$P\left(\bar{\mathbf{X}}^0(j+1) = (y_1, y_2) | \bar{\mathbf{X}}^0(j) = (x_1, x_2)\right) = \begin{cases} \frac{\Gamma^1_{x_1,y_1}}{q_1(x_1)+q_2(x_2)}, & y_1 \neq x_1,\ y_2 = x_2, \\ \frac{\Gamma^2_{x_2,y_2}}{q_1(x_1)+q_2(x_2)}, & y_1 = x_1,\ y_2 \neq x_2, \\ 0, & \text{otherwise.} \end{cases}$$

The dynamics thus describe a process for which, when the current state is (x_1, x_2), the support of the jump distribution is $\{(y_1, x_2) : y_1 \in \mathcal{S}\} \cup \{(x_1, y_2) : y_2 \in \mathcal{S}\}$. The jump times of $\mathbf{X}^0$ are exponential random variables with jump rates given by $q(x_1, x_2) = q_1(x_1) + q_2(x_2)$. This is conveniently summarized by the infinitesimal generator $\mathcal{L}^0$ associated with the process $\mathbf{X}^0$:

$$\begin{aligned} \mathcal{L}^0 f(x_1, x_2) = &\frac{1}{q(x_1, x_2)} \sum_{y_1 \neq x_1} \left[f(y_1, x_2) - f(x_1, x_2)\right] \Gamma^1_{x_1,y_1} \\ &+ \frac{1}{q(x_1, x_2)} \sum_{y_2 \neq x_2} \left[f(x_1, y_2) - f(x_1, x_2)\right] \Gamma^2_{x_2,y_2}. \end{aligned}$$

Let η_T^0 denote the empirical measure of $\mathbf{X}^0$,

$$\eta_T^0(\cdot) = \frac{1}{T} \int_0^T \delta_{\mathbf{X}^0(t)}(\cdot) dt.$$

By ergodicity, with probability one η_T^0 converges weakly to μ in $\mathcal{P}(\mathcal{S}^2)$, the space of probability measures on $\mathcal{S}^2$, as $T \to \infty$ and the empirical measure η_T^0 is used to approximate ergodic averages associated with μ. In particular, this provides a way to obtain estimates of integrals associated with μ_1, the invariant measure associated with the lower temperature and often the distribution of interest in practice. However, the problem of rare-event sampling may cause the rate of convergence for η_T^0 to μ to be slow, resulting in inaccurate estimates, especially with respect to the low temperature marginal.

Let $\mathbf{X}^a$ be the process that corresponds to swaps according to a Metropolis-type rule and with rate $a \geq 0$. More precisely, let b be the function

$$b(x_1, x_2) = 1 \wedge \frac{\mu(x_2, x_1)}{\mu(x_1, x_2)} \tag{2.1}$$

and let the infinitesimal generator $\mathcal{L}^a$ associated with $\mathbf{X}^a$ be given by

$$\mathcal{L}^a f(x_1, x_2) = \mathcal{L}^0 f(x_1, x_2) + ab(x_1, x_2)\left[f(x_2, x_1) - f(x_1, x_2)\right].$$

$\mathbf{X}^a$ is a continuous time version of the well-known parallel tempering algorithm [7, 11, 16, 17].

The process $\mathbf{X}^a$ has two different kinds of jumps. Suppose the process is in (x_1, x_2) and let s_1 and s_2 be two exponential random variables with parameter $q(x_1, x_2)$ and a, respectively. Then the process has a jump after $s = s_1 \wedge s_2$ units of time. Note that the time s to the next jump is an exponential random variable with parameter $q(x_1, x_2) + a$. If $s_1 < s_2$ the jump is according to the dynamics specified by Γ^1 and Γ^2 (the $\mathcal{L}^0$ part of the generator). If instead $s_2 < s_1$, with probability $b(x_1, x_2)$ the process jumps to (x_2, x_1) - the two particles switch locations - and with remaining probability it stays in (x_1, x_2) (a failed swap attempt). In terms of the embedded Markov chain $\bar{\mathbf{X}}^a$ this corresponds to a probability transition kernel

$$\begin{aligned}
&P\left(\bar{\mathbf{X}}^a(j+1) = (y_1, y_2) | \bar{\mathbf{X}}^a(j) = (x_1, x_2)\right) \\
&= \frac{q(x_1, x_2)}{a + q(x_1, x_2)} P\left(\bar{\mathbf{X}}^0(j+1) = (y_1, y_2) | \bar{\mathbf{X}}^0(j) = (x_1, x_2)\right) \\
&\quad + \frac{a}{a + q(x_1, x_2)}\left[b(x_1, x_2)\delta_{(x_2, x_1)}(y_1, y_2) + (1 - b(x_1, x_2))\delta_{(x_1, x_2)}(y_1, y_2)\right].
\end{aligned}$$

As previously mentioned the jump times of $\mathbf{X}^a$ are exponential random variables with jump rates $q(x_1, x_2) + a$ (when in state (x_1, x_2)). Although easy to check using detailed balance, it is important to note that introducing the

second type of jump - swaps of the particle locations - does not change the invariant distribution; just as for $\mathbf{X}^0$, μ is the invariant distribution of $\mathbf{X}^a$.

The infinite swapping process, first studied in [6], is the limit process that arises as $a \to \infty$. However, the processes $\{\mathbf{X}^a\}$ cannot be tight as a grows due to the discontinuities introduced by the swapping of particle locations and there is no well-defined limit process. Instead, a process associated with the limit $a \to \infty$ can be obtained by considering a temperature swapped version of $\mathbf{X}^a$, denoted $\mathbf{Y}^a$. We give only a brief description of $\mathbf{Y}^a$ and the reader is referred to [6] for details and a more thorough discussion.

Rather than attempting to swap particle locations, one can let dynamics of the different components of the Markov process $\mathbf{X}^0$ change. To describe this, suppose we append a process Z^a with values in $\{1, 2\}$, and consider the Markov process $(\mathbf{Y}^a, Z^a)$ with jump rates given $(\mathbf{Y}^a, Z^a) = (x_1, x_2, z)$ equal to

$$\begin{array}{ll} \Gamma^1_{x_1,y_1} & \text{for } (x_1, x_2, 1) \to (y_1, x_2, 1) \\ \Gamma^2_{x_2,y_2} & \text{for } (x_1, x_2, 1) \to (x_1, y_2, 1) \\ \Gamma^2_{x_1,y_1} & \text{for } (x_1, x_2, 2) \to (y_1, x_2, 2) \\ \Gamma^1_{x_2,y_2} & \text{for } (x_1, x_2, 2) \to (x_1, y_2, 2) \\ ab(x_1, x_2) & \text{for } (x_1, x_2, 1) \to (x_1, x_2, 2) \\ ab(x_2, x_1) & \text{for } (x_1, x_2, 2) \to (x_1, x_2, 1) \end{array} .$$

With this process the particles do not change location when a swap is successful. Instead, the dynamics are swapped as indicated by the value of Z^a. To account for this change the empirical measure must be redefined, and in fact one uses

$$\frac{1}{T}\int_0^T \left[1_{\{Z^a(t)=1\}}\delta_{(\mathbf{Y}^a_1(t),\mathbf{Y}^a_2(t))}(\cdot) + 1_{\{Z^a(t)=2\}}\delta_{(\mathbf{Y}^a_2(t),\mathbf{Y}^a_1(t))}(\cdot)\right] dt \tag{2.2}$$

in lieu of

$$\frac{1}{T}\int_0^T \delta_{\mathbf{X}^a(t)}(\cdot)dt.$$

Because of the change in bookkeeping, the process $\mathbf{Y}^a$ as well as the replacement for the empirical measure have well defined limits in distribution as $a \to \infty$ [6]. Let $\mathbf{Y}^\infty = \{(Y^\infty_1(t), Y^\infty_2(t)) : t \geq 0\}$ denote the limit process, referred to as the infinite swapping limit or the infinite swapping process. It follows from the dynamics of $\mathbf{Y}^a$ that $\mathbf{Y}^\infty$ is a pure jump Markov process with infinitesimal generator

$$\mathcal{L}^\infty f(x_1, x_2) = \sum_{(y_1,y_2)\in\mathcal{S}^2} [f(y_1, y_2) - f(x_1, x_2)]\, \Gamma^\infty_{(x_1,x_2),(y_1,y_2)} \tag{2.3}$$

where the rate matrix Γ^∞ is defined as

$$\Gamma^\infty_{(x_1,x_2),(y_1,y_2)} = \begin{cases} \rho(x_1,x_2)\Gamma^1_{x_1,y_1} + \rho(x_2,x_1)\Gamma^2_{x_1,y_1}, & y_1 \neq x_1, y_2 = x_2, \\ \rho(x_1,x_2)\Gamma^2_{x_2,y_2} + \rho(x_2,x_1)\Gamma^1_{x_2,y_2}, & y_1 = x_1, y_2 \neq x_2, \\ 0, & \text{otherwise}, \end{cases} \tag{2.4}$$

where

$$\rho(x_1,x_2) = \frac{\mu(x_1,x_2)}{\mu(x_1,x_2)+\mu(x_2,x_1)} \tag{2.5}$$

is the relative weight μ assigns to the permutation (x_1,x_2). The limit of (2.2) in distribution is

$$\frac{1}{T}\int_0^T \left[\rho(\mathbf{Y}_1^\infty(t),\mathbf{Y}_2^\infty(t))\delta_{(\mathbf{Y}_1^\infty(t),\mathbf{Y}_2^\infty(t))}(\cdot) + \rho(\mathbf{Y}_2^\infty(t),\mathbf{Y}_1^\infty(t))\delta_{(\mathbf{Y}_2^\infty(t),\mathbf{Y}_1^\infty(t))}(\cdot)\right] dt. \tag{2.6}$$

In the general case, $K \geq 2$, ρ is defined similarly, with the denominator in (2.5) now being the sum over of $\mu(\mathbf{x}^\sigma)$ for all permutations $\sigma \in \Sigma_K$, and the replacement for the empirical measure using a sum over all such permutations. It is shown in [6] that the infinite swapping approximation (2.2) converges to μ faster than the corresponding parallel tempering scheme (i.e., the empirical measure of $\mathbf{X}^a$) for any $a \in [0,\infty)$, when the corresponding large deviations rate functions are used to measure the rate of convergence.

Similar to the pre-limit dynamics, $\Gamma^\infty_{(x_1,x_2),(x_1,x_2)} = -\sum_{(y_1,y_2)\in\mathcal{S}^2}\Gamma^\infty_{(x_1,x_2),(y_1,y_2)}$ and we define $q^\infty(x_1,x_2) = -\Gamma^\infty_{(x_1,x_2),(x_1,x_2)}$. The interpretation is that the total rate out of state (x_1,x_2) for the infinite swapping process is a mixture, according to the weights $\rho(x_1,x_2)$ and $\rho(x_2,x_1)$, of the rates out of (x_1,x_2) and (x_2,x_1) for the original (uncoupled) process. The following symmetry properties will come in handy later on. Although not obvious, they follow immediately from the definitions of Γ^∞ and q^∞ and the proof is merely a matter of bookkeeping.

Lemma 2.1 *For any $\mathbf{x},\mathbf{y} \in \mathcal{S}^K$ and permutation $\sigma \in \Sigma_K$, $q^\infty(\mathbf{x}) = q^\infty(\mathbf{x}^\sigma)$ and $\Gamma^\infty_{\mathbf{x},\mathbf{y}} = \Gamma^\infty_{\mathbf{x}^\sigma,\mathbf{y}^\sigma}$.*

Proof. For notational simplicity we limit the proof to the case $K = 2$. Start by considering the claim for Γ^∞. If (x_1,x_2) and (y_1,y_2) are chosen such that $\Gamma^\infty_{(x_1,x_2),(y_1,y_2)} = 0$ then $\Gamma^\infty_{(x_2,x_1),(y_2,y_1)} = 0$ as well and the claim holds. Without loss of generality we can consider $y_1 \neq x_1$ and $y_2 = x_2$, so that

$$\Gamma^\infty_{(x_1,x_2),(y_1,y_2)} = \rho(x_1,x_2)\Gamma^1_{x_1,y_1} + \rho(x_2,x_1)\Gamma^2_{x_1,y_1}.$$

Labeling the states $(u_1, u_2) = (x_2, x_1)$ and $(v_1, v_2) = (y_2, y_1)$, we have $v_1 = u_1$ and $v_2 \neq u_2$. By (2.4),

$$\begin{aligned}\Gamma^{\infty}_{(x_2,x_1),(y_2,y_1)} &= \Gamma^{\infty}_{(u_1,u_2),(v_1,v_2)} \\ &= \rho(u_1,u_2)\Gamma^{2}_{u_2,v_2} + \rho(u_2,u_1)\Gamma^{1}_{u_2,v_2} \\ &= \rho(x_2,x_1)\Gamma^{2}_{x_1,y_1} + \rho(x_1,x_2)\Gamma^{1}_{x_1,y_1} \\ &= \Gamma^{\infty}_{(x_1,x_2),(y_1,y_2)}.\end{aligned}$$

The case where $y_1 = x_1, y_2 \neq x_2$ is completely analogous. This confirms the claim for Γ^{∞} and the symmetry property for q^{∞} then follows directly from the definition:

$$\begin{aligned}q^{\infty}(x_1,x_2) &= \sum_{(y_1,y_2)\in\mathcal{S}^2} \Gamma^{\infty}_{(x_1,x_2),(y_1,y_2)} \\ &= \sum_{(y_1,y_2)\in\mathcal{S}^2} \Gamma^{\infty}_{(x_2,x_1),(y_2,y_1)} \\ &= q^{\infty}(x_2,x_1).\end{aligned}$$

The extension to $K > 2$ is straightforward. ■

As is discussed at the beginning of Section 3, standard arguments using detailed balance show that the stationary distribution for $\mathbf{Y}^{\infty}$ is $\bar{\mu}$, a symmetrized version of μ:

$$\bar{\mu}(y_1,y_2) = \frac{1}{2}[\mu(y_1,y_2) + \mu(y_2,y_1)].$$

It follows that the weighted empirical measure

$$\eta^{\infty}_{T} = \frac{1}{T}\int_0^T \left[\rho(Y_1^{\infty}(t),Y_2^{\infty}(t))\delta_{(Y_1^{\infty}(t),Y_2^{\infty}(t))} + \rho(Y_2^{\infty}(t),Y_1^{\infty}(t))\delta_{(Y_2^{\infty}(t),Y_1^{\infty}(t))}\right] dt \tag{2.7}$$

converges to μ as $T \to \infty$: using (2.5), for any test function f

$$\begin{aligned}&\sum_{(y_1,y_2)\in S^2} \left[f(y_1,y_2)\rho(y_1,y_2) + f(y_2,y_1)\rho(y_2,y_1)\right]\bar{\mu}(y_1,y_2) \qquad (2.8)\\ &= 2\sum_{(y_1,y_2)\in S^2} f(y_1,y_2)\rho(y_1,y_2)\frac{1}{2}\left[\mu(y_1,y_2)+\mu(y_2,y_1)\right] \\ &= \sum_{(y_1,y_2)\in S^2} f(y_1,y_2)\mu(y_1,y_2).\end{aligned}$$

Thus, η_T^∞ can be used as an approximation of μ and for computing ergodic averages of thermodynamic properties of the original process.

In the case of K temperatures, there are K rate matrices Γ^i. Let μ denote the product measure of the associated μ_i's. The analogue of the claim just made still holds and η_T^∞ takes the form

$$\eta_T^\infty = \frac{1}{T}\int_0^T \sum_{\sigma\in\Sigma_K} \rho((\mathbf{Y}^\infty(t))^\sigma)\delta_{(\mathbf{Y}^\infty(t))^\sigma}\,dt,$$

where Σ_K is the set of permutations of $\{1,2,\ldots,K\}$ and $\mathbf{Y}^\infty$ is now a Markov process that has generator (2.3) and with Γ^∞ defined accordingly.

For notational simplicity, in general we do not distinguish between two or more temperatures. Whenever a proof is provided only for $K = 2$, unless otherwise stated, the reader should convince themselves that extending the result to an arbitrary number of temperatures is a straightforward task.

2.2 Particle-temperature associations

As mentioned in the Introduction, when running parallel tempering or infinite swapping algorithms, in addition to the empirical measure used for computing ergodic averages one can record the empirical measure on the particle-temperature associations, i.e., the fraction of time that a given particle in the $\mathbf{Y}^a$ or $\mathbf{Y}^\infty$ formulation is assigned the dynamic Γ^i.

To discuss particle-temperature associations it is useful to consider first the pre-limit process $\mathbf{Y}^a$. Recall from the previous subsection that $\mathbf{Y}^a$ is the process for which the dynamics, i.e., the rate matrices, associated with each particle (component of $\mathbf{Y}^a$) are changed according to swaps attempted with intensity a, and with the probability of success for each attempt given by b in (2.1). We think of each permutation as a mapping from $\{1,2,\ldots,K\}$ onto itself. One can imagine starting with σ equal to the identity permutation, and updating σ each time a swap of the dynamics occurs. Thus at each moment of time there is a particular permutation, σ, that provides the current assignment of rate matrices to the components $Y_1^a,\ldots,Y_K^a$: $\Gamma^{\sigma^{-1}(i)}$ is the rate matrix associated with particle Y_i^a.

For the pre-limit processes, the particle-temperature associations are defined as the fraction of time that a certain permutation $\sigma\in\Sigma_K$ is used to associate rate matrices to the components of the process. This is then used to create an empirical measure, $\rho_T = (\rho_T^{\sigma_1},\ldots,\rho_T^{\sigma_{K!}})$ (for some ordering of the elements of Σ_K) on the set Σ_K.

Although the possibility of recording what rate matrix is associated with what component is blurred in the infinite swapping limit, the definition of

the particle-temperature associations still make sense: For each permutation $\sigma \in \Sigma_K$, the corresponding component ρ_T^σ of ρ_T is defined by

$$\rho_T^\sigma = \frac{1}{T}\int_0^T \rho((\mathbf{Y}^\infty(t))^\sigma)dt,\ \rho_T = \{\rho_T^\sigma\}_{\sigma\in\Sigma_K}.$$

Thus ρ_T is probability measure that puts mass ρ_T^σ on the permutation σ. In the case of two temperatures we typically denote the components by $\rho_T = (\rho_T^1, \rho_T^2)$, where the superscript 1 refers to the identity map $\sigma = (1, 2)$ and 2 to the permutation that reverses components, $\sigma = (2, 1)$. Note that ρ_T is a vector with $|\Sigma_K| = K!$ entries, each providing the fraction of time the corresponding permutation has been in use up to time T. Since there must always be some permutation that is in use this vector is a probability measure on Σ_K.

To further understand the role of ρ_T consider again $\mathbf{Y}^a$, the process associated with parallel tempering with swap rate a (and with swapped dynamics). It is not hard to show that, if the state of $\mathbf{Y}^a$ is frozen at some $\mathbf{y}$, then the swap mechanism introduces an ergodic Markov chain on Σ_K, with $\rho(\mathbf{y}^\sigma)$ the stationary probability to be in permutation σ. Indeed, in the case $K = 2$, if we label the permutations σ_1 and σ_2, a successful swap from σ_1 to σ_2 has probability $1 \wedge (\mu(y_2, y_1)/\mu(y_1, y_2))$ and a successful swap in the reverse direction has probability $1 \wedge (\mu(y_1, y_2)/\mu(y_2, y_1))$. Without loss of generality we can assume that the former probability is $\mu(y_2, y_1)/\mu(y_1, y_2)$, so that the probability of a successful swap from σ_2 to σ_1 is 1. It is clear that these transitions form a Markov chain on $\Sigma_2 = \{\sigma_1, \sigma_2\}$ and under the assumption on the swap probabilities the transition matrix is

$$\begin{pmatrix} 1 - \frac{\mu(y_2,y_1)}{\mu(y_1,y_2)} & \frac{\mu(y_2,y_1)}{\mu(y_1,y_2)} \\ 1 & 0 \end{pmatrix}.$$

It is easy to check that the associated invariant measure is that which puts probability $\rho(y_1, y_2)$ on permutation σ_1 and $\rho(y_2, y_1)$ on σ_2. The generalization to arbitrary $K \geq 2$ is straightforward, albeit notationally cumbersome.

Thus for a fixed $\mathbf{y}$, $\rho(\mathbf{y}^\sigma)$ can be interpreted as the asymptotic fraction of time that temperatures are assigned according to σ when $\mathbf{Y}^\infty(t) = \mathbf{y}$.

Under the invariant distribution $\bar{\mu}$ of $\mathbf{Y}^\infty$ the expectation of $\rho((\mathbf{Y}^\infty(t))^\sigma)$ is

$$\begin{aligned} E\left[\rho((\mathbf{Y}^\infty(t))^\sigma)\right] &= \sum_{\mathbf{y}\in\mathcal{S}^K} \rho(y_{\sigma^{-1}(1)},\ldots,y_{\sigma^{-1}(K)})\bar{\mu}(y_1,\ldots,y_K) \\ &= \sum_{\mathbf{y}\in\mathcal{S}^K} \left(\frac{\mu(y_{\sigma^{-1}(1)},\ldots,y_{\sigma^{-1}(K)})}{\sum_{\bar{\sigma}\in\Sigma_K}\mu(y_{\bar{\sigma}^{-1}(1)},\ldots,y_{\bar{\sigma}^{-1}(K)})}\right) \\ &\qquad \times \frac{1}{K!}\frac{\sum_{\bar{\sigma}\in\Sigma_K}\mu(y_{\bar{\sigma}^{-1}(1)},\ldots,y_{\bar{\sigma}^{-1}(K)})}{\prod_{k=1}^K Z_k} \\ &= \frac{1}{K!}\sum_{\mathbf{y}\in\mathcal{S}^K}\frac{\mu(y_{\sigma^{-1}(1)},\ldots,y_{\sigma^{-1}(K)})}{\prod_{k=1}^K Z_k} \\ &= \frac{1}{K!}. \end{aligned}$$

That is, under the invariant distribution $\bar{\mu}$ all permutations are assigned the same probability. It follows from the ergodic theorem that ρ_T converges to the uniform distribution on Σ_K.

3 Large deviation properties

To use ρ_T (or related quantities) to understand η_T^∞, we need to study the joint distribution of η_T^∞ and ρ_T as T grows to infinity. More precisely, we will study the asymptotic properties of this joint distribution by means of large deviations. In [6] the large deviation properties associated with infinite swapping are explored by considering the measure η_T^∞ in the limit as T goes to infinity. Here, the starting point is instead the ordinary empirical measure associated with the infinite swapping process $\mathbf{Y}^\infty$:

$$\nu_T(\cdot) \doteq \frac{1}{T}\int_0^T \delta_{\mathbf{Y}^\infty(t)}(\cdot)dt.$$

As noted in the previous section, $\mathbf{Y}^\infty$ is a continuous time pure jump process with generator given in (2.3). It is assumed that the rate matrices Γ^1, Γ^2 are reversible with respect to μ_1 and μ_2, respectively. Under this assumption it is straightforward to show that Γ^∞ is reversible with respect to $\bar{\mu}$. Without loss of generality, for some $\mathbf{x} = (x_1, x_2) \in \mathcal{S}^2$ take $\mathbf{y} = (y_1, x_2)$, so that

$$\Gamma^\infty_{\mathbf{x},\mathbf{y}} = \rho(x_1, x_2)\Gamma^1_{x_1,y_1} + \rho(x_2, x_1)\Gamma^2_{x_1,y_1}.$$

It then follows that

$$\begin{aligned}\bar{\mu}(\mathbf{x})\Gamma^{\infty}_{\mathbf{x},\mathbf{y}} &= \frac{1}{2}\left[\mu((x_1,x_2)) + \mu((x_2,x_1))\right]\left(\rho(x_1,x_2)\Gamma^1_{x_1,y_1} + \rho(x_2,x_1)\Gamma^2_{x_1,y_1}\right)\\ &= \frac{1}{2}\mu((x_1,x_2))\Gamma^1_{x_1,y_1} + \frac{1}{2}\mu((x_2,x_1))\Gamma^2_{x_1,y_1}\\ &= \frac{1}{2}\mu((y_1,x_2))\Gamma^1_{y_1,x_1} + \frac{1}{2}\mu((x_2,y_1))\Gamma^2_{y_1,x_1}\\ &= \bar{\mu}((y_1,x_1))\left(\rho(y_1,x_2)\Gamma^1_{y_1,x_1} + \rho(x_2,y_1)\Gamma^2_{y_1,x_1}\right)\\ &= \bar{\mu}(\mathbf{y})\Gamma^{\infty}_{\mathbf{y},\mathbf{x}},\end{aligned}$$

where we have used the definitions of ρ and Γ^{∞}, $\mu(x_1,x_2) = \mu_1(x_1)\mu_2(x_2)$, and the reversibility of the rate matrices with respect to μ_1 and μ_2. For other choices of $\mathbf{y}$ the calculations are completely analogous.

In addition to reversibility, for the purpose of large deviation results we assume that μ_1 and μ_2 are the unique invariant measures of the Markov processes with rate matrices Γ^1 and Γ^2, respectively. In the finite state setting considered here, uniqueness of the invariant distributions can be ensured by assuming that all states communicate. Note that under this condition $\mu(\mathbf{x}) > 0$ for all $\mathbf{x} \in \mathcal{S}^K$. For a general compact state space $\mathcal{S}$ see [5] for a set of sufficient conditions for the large deviation results to hold.

By the results of [5], the sequence $\{\nu_T\}_T$ satisfies a large deviation principle on $\mathcal{P}(\mathcal{S}^K)$ with convex and lower semicontinuous rate function J given by

$$J(\nu) = \sum_{\mathbf{x}\in\mathcal{S}^K} q^{\infty}(\mathbf{x})\theta(\mathbf{x})\bar{\mu}(\mathbf{x}) - \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^K} \theta^{1/2}(\mathbf{x})\theta^{1/2}(\mathbf{y})\Gamma^{\infty}_{\mathbf{x},\mathbf{y}}\bar{\mu}(\mathbf{x}),$$

where $\theta(\mathbf{x}) = \nu(\mathbf{x})/\bar{\mu}(\mathbf{x})$, $\mathbf{x} \in \mathcal{S}^K$. We show in Lemma A.3 in the appendix that J is strictly convex.

The empirical measure ν^{∞}_T is of interest because although it is η^{∞}_T that is used for all computational purposes, in an infinite swapping algorithm it is ν^{∞}_T that is simulated and from which one obtains η^{∞}_T. The explicit connection between the two empirical measures is through the mapping $M : \mathcal{P}(\mathcal{S}^K) \to \mathcal{P}(\mathcal{S}^K)$ given by

$$(M\nu)(\mathbf{x}) = \rho(x_1,\ldots,x_K)\sum_{\sigma\in\Sigma_K}\nu(x_{\sigma^{-1}(1)},\ldots,x_{\sigma^{-1}(K)}), \tag{3.1}$$

for $\mathbf{x} \in \mathcal{S}^K$, $\nu \in \mathcal{P}(\mathcal{S}^K)$. For $K = 2$ the definition simplifies to

$$(M\nu)(x_1,x_2) = \rho(x_1,x_2)[\nu(x_1,x_2) + \nu(x_2,x_1)]. \tag{3.2}$$

To see that (3.1) maps probability measures to probability measures, take any $\nu \in \mathcal{P}(\mathcal{S}^K)$ [see (2.8)]. Then since $\sum_{\sigma\in\Sigma_K} \rho^\sigma(\mathbf{x}) = 1$ for any $\mathbf{x} \in \mathcal{S}$,

$$\begin{aligned}
\sum_{\mathbf{x}\in\mathcal{S}^K} (M\nu)(\mathbf{x}) &= \sum_{\mathbf{x}\in\mathcal{S}^K} \left(\rho(\mathbf{x}) \sum_{\sigma\in\Sigma_K} \nu(x_{\sigma^{-1}(1)}, \ldots, x_{\sigma^{-1}(K)}) \right) \\
&= \sum_{\sigma\in\Sigma_K} \sum_{\mathbf{x}\in\mathcal{S}^K} \rho(\mathbf{x})\nu(\mathbf{x}^\sigma) \\
&= \sum_{\sigma\in\Sigma_K} \sum_{\mathbf{x}\in\mathcal{S}^K} \rho(\mathbf{x}^\sigma)\nu(\mathbf{x}) \\
&= \sum_{\mathbf{x}\in\mathcal{S}^K} \sum_{\sigma\in\Sigma_K} \rho(\mathbf{x}^\sigma)\nu(\mathbf{x}) \\
&= 1.
\end{aligned}$$

Thus $M\nu$ is indeed a probability measure on $\mathcal{S}^K$. Furthermore, it is easy to check that $\eta_T^\infty = M\nu_T$ and $\mu = M\bar{\mu}$. In general, M is the map that takes a measure associated with the symmetrized variables and maps it to a corresponding measure for the unsymmetrized variables.

The form of M for the case of two temperatures, given in (3.2), highlights a sort of symmetric property of the mapping: As long as the total mass under ν of the two points (x, y) and (y, x) is kept fixed, the image measure $M\nu$ will remain the same. The analogue of this of course holds for $K > 2$ as well and it turns out to be an essential property for studying infinite swapping.

Once the mapping M that takes ν_T to η_T^∞ has been identified, together with the large deviations principle for ν_T, the joint large deviation principle for (η_T^∞, ρ_T) is obtained through an application of the contraction principle. For a function f and probability measure γ on $\mathcal{S}^K$ let $\langle f, \gamma\rangle = \sum_{\mathbf{x}\in\mathcal{S}^K} f(\mathbf{x})\gamma(\mathbf{x})$.

Proposition 3.1 *Suppose that μ_1 and μ_2 are the unique invariant measures for Γ^1 and Γ^2, respectively. The sequence $\{(\eta_T^\infty, \rho_T)\}$ satisfies a large deviation principle on $\mathcal{P}(\mathcal{S}^K) \times \mathcal{P}(\Sigma_K)$ with rate function*

$$I(\gamma, w) = \inf \{ J(\nu) : \ \gamma = M\nu, \ \langle \rho^\sigma, \nu\rangle = w_\sigma, \ \sigma \in \Sigma_K \} .$$

Proof. The map M is continuous with respect to the weak topology (which is the same as the standard Euclidean topology if we think of $\mathcal{P}(\mathcal{S}^K)$ as embedded into $\mathbb{R}^{|\mathcal{S}|^K}$). Furthermore, the components of ρ_T can be expressed as expectations with respect to ν_T:

$$\rho_T^\sigma = \sum_{\mathbf{x}\in\mathcal{S}^K} \rho(\mathbf{x}^\sigma)\nu_T(\mathbf{x}) = \langle \rho^\sigma, \nu_T\rangle, \ \sigma \in \Sigma_K.$$

Consider the map from $\mathcal{P}(\mathcal{S}^K)$ to $\mathcal{P}(\Sigma_K)$ defined by $\nu \mapsto (\langle \rho^{\sigma_1}, \nu \rangle, \ldots, \langle \rho^{\sigma_{K!}}, \nu \rangle)$. This is a continuous map with respect to the weak topology and precisely the one that takes ν_T to ρ_T. Thus, the pair (η_T^∞, ρ_T) is obtained by applying a continuous map to ν_T. It follows from the contraction principle that the sequence satisfies a large deviation principle with the prescribed rate function as $T \to \infty$. ■

In this paper we start with the symmetrized process $\mathbf{Y}^\infty$ and its associated empirical measure, and then study various large deviation properties of infinite swapping through mappings and the contraction principle. In [6] the large deviation principle for η_T^∞ is proved without reference to the symmetrized process (though for diffusion processes rather than the jump processes discussed here). For completeness some comments on the relation between the rate directly obtained in [6] (more precisely the analogous rate appropriate for the model considered here) and the joint large deviation principle of Proposition 3.1 are appropriate.

Note that from the joint large deviation principle one immediately obtains large deviation principles for the marginals η_T^∞ and ρ_T. Denote the corresponding rate functions I_1 and I_2:

$$I_1(\gamma) = \inf \{ J(\nu) : \ M\nu = \gamma \}, \ \gamma \in \mathcal{P}(\mathcal{S}^K),$$

and

$$I_2(w) = \inf \{ J(\nu) : \ \langle \rho^\sigma, \nu \rangle = w_\sigma, \ \sigma \in \Sigma_K \}, \ w = \{w_\sigma\} \in \mathcal{P}(\Sigma_K). \quad (3.3)$$

Let I^∞ denote the rate function analogous to that of [6] that would be appropriate for $\{\eta_T^\infty\}$,

$$I^\infty(\gamma) = \sum_{\mathbf{x} \in \mathcal{S}^2} q(\mathbf{x}) \gamma(\mathbf{x}) - \sum_{\mathbf{x},\mathbf{y} \in \mathcal{S}^2} \left[\frac{d\gamma}{d\mu} \right]^{1/2} (\mathbf{x}) \left[\frac{d\gamma}{d\mu} \right]^{1/2} (\mathbf{y}) \Gamma_{\mathbf{x},\mathbf{y}} \mu(\mathbf{x})$$

where Γ is the rate matrix associated with the original uncoupled processes. This rate function is finite only for those measures γ that satisfy, for $K = 2$,

$$[d\gamma/d\mu](x_1, x_2) = [d\gamma/d\mu](x_2, x_1). \quad (3.4)$$

This constraint is immediately satisfied by any measure γ on the form $\gamma = M\nu$, for some $\nu \in \mathcal{P}(\mathcal{S}^2)$, and thus it is built in to the definitions of I and I_1. That is, if γ violates this condition of correct relative weights on (x_1, x_2) and (x_2, x_1), then there is no ν for which $\gamma = M\nu$ and both I and I_1 are by definition infinite for such measures. To see that this is true is a short

calculation using only the definition of M: Take $\gamma \in \mathcal{P}(\mathcal{S}^2)$ for which there is some ν such that $\gamma = M\nu$ and $\gamma \ll \mu$. Since these are discrete measures, for any $(x_1, x_2) \in \mathcal{S}^2$,

$$\begin{aligned}\left[\frac{d\gamma}{d\mu}\right](x_1, x_2) &= \frac{\gamma(x_1, x_2)}{\mu(x_1, x_2)} \\ &= \frac{(M\nu)(x_1, x_2)}{\mu(x_1, x_2)} \\ &= \frac{\rho(x_1, x_2)[\nu(x_1, x_2) + \nu(x_2, x_1)]}{\mu(x_1, x_2)} \\ &= \frac{\nu(x_1, x_2) + \nu(x_2, x_1)}{2\bar{\mu}(x_1, x_2)},\end{aligned}$$

and it follows by symmetry that $[d\gamma/d\mu](x_2, x_1)$ is equal to this as well. It should be clear that the relation also holds for $K > 2$. Furthermore, it continues to be true in the case of an uncountable state space $\mathcal{S}$, most easily seen by assuming that all measures involved have densities with respect to some common reference measure and using the same kind of argument as here.

Although the INS dynamics use symmetrized dynamics, there is no reason that its empirical measure must be symmetric. However, if γ is a point in the support of the weighted empirical measure (2.7), then the most likely empirical measure for the INS process that leads to γ will be symmetric. Thus for probability measures γ that satisfy the weighted symmetry condition (3.4) we have the following result.

Proposition 3.2 *Suppose $\gamma \in \mathcal{P}(\mathcal{S}^K)$ is such that $\mathcal{M}(\gamma) = \{\nu \in \mathcal{P}(\mathcal{S}^K) : \gamma = M\nu\}$ is non-empty. Then*

$$\inf\{J(\nu) : \nu \in \mathcal{M}(\gamma)\}$$

is attained at the symmetric ν_{sym} given by

$$\nu_{sym}(x_1, \ldots, x_K) = \frac{\gamma(x_1, \ldots, x_K)}{K!\rho(\mathbf{x})}. \tag{3.5}$$

Before proceeding with the proof a short remark on the form of ν_{sym} is be in place. Although not necessarily apparent at first, ν_{sym} is indeed symmetric:

$$\nu_{sym}(x_1, \ldots, x_K) = \nu_{sym}(x_{\sigma^{-1}(1)}, \ldots, x_{\sigma^{-1}(K)}),$$

for any $\mathbf{x} \in \mathcal{S}^K$ and permutation $\sigma \in \Sigma_K$. This follows from the weighted symmetry of γ as expressed in (3.4) which holds for any measure obtained as a mapping through M, and the definition of ρ. For the sake of clarity we show for $K = 2$ that under the mapping M, ν_{sym} returns the measure γ. Indeed, for any $(x_1, x_2) \in \mathcal{S}^2$,

$$\begin{aligned}
(M\nu_{sym})(x_1, x_2) &= \rho(x_1, x_2)\left[\nu_{sym}(x_1, x_2) + \nu_{sym}(x_2, x_1)\right] \\
&= \rho(x_1, x_2)\left(\frac{\gamma(x_1, x_2)}{2\rho(x_1, x_2)} + \frac{\gamma(x_2, x_1)}{2\rho(x_2, x_1)}\right) \\
&= \frac{\mu(x_1, x_2)}{2}\left(\frac{\gamma(x_1, x_2)}{\mu(x_1, x_2)} + \frac{\gamma(x_2, x_1)}{\mu(x_2, x_1)}\right).
\end{aligned}$$

The condition (3.4) on the relative weights $[d\gamma/d\mu]$ then shows that this is equal to $\gamma(x_1, x_2)$.

Since the rate function of a large deviation principle is unique the following result is to be expected.

Corollary 3.3 *For any $\gamma \in \mathcal{P}(\mathcal{S}^K)$, $I_1(\gamma) = I^\infty(\gamma)$.*

The result follows almost immediately from Proposition 3.2 by inserting the symmetric measure ν_{sym} into J, using the definitions of γ^∞ and Γ^∞ and the constraint (3.4). The calculations are straightforward but cumbersome and are left out for brevity.

Proof of Proposition 3.2. Due to the form of M, $\mathcal{M}(\gamma)$ is a closed, convex set in $\mathcal{P}(\mathcal{S}^K)$. Thus, if $\mathcal{M}(\gamma)$ is non-empty, then by strict convexity J will uniquely attain its infimum over the set. Using this fact we argue by convexity that the minimizing measure must indeed be the symmetric ν_{sym}.

For any $\nu \in \mathcal{M}(\gamma)$ let ν^σ be a permutation of ν according to σ. For any set $A \subset \mathcal{S}^K$,

$$\nu^\sigma(A) = \sum_{\mathbf{x} \in \mathcal{S}^K} I\{(x_1, \ldots, x_K) \in A\}\nu(d\mathbf{x}^\sigma),$$

or, equivalently, $\nu^\sigma(A) = \nu(A^\sigma)$, where $A^\sigma = \{\mathbf{x}^\sigma : \mathbf{x} \in A\}$. The symmetry property of M, discussed after (3.2), ensures that $\nu^\sigma \in \mathcal{M}(\gamma)$ for every $\sigma \in \Sigma_K$. Moreover, since the measure $\bar{\mu}$ is symmetric, $\nu \ll \bar{\mu}$ implies $\nu^\sigma \ll \bar{\mu}$. The key to the proof is to show that $J(\nu) = J(\nu^\sigma)$. Indeed, assume that this holds and that $\bar{\nu}$ is the unique minimizing measure in $\mathcal{M}(\gamma)$,

$$\bar{\nu} = \operatorname*{argmin}_{\nu \in \mathcal{M}(\gamma)} J(\nu).$$

If $\bar{\nu}$ is not symmetric, then there is a $\sigma \in \Sigma_K$ for which $\bar{\nu} \neq \bar{\nu}^\sigma$ and it follows that $\bar{\nu} \neq (1/2)\bar{\nu} + (1/2)\bar{\nu}^\sigma \in \mathcal{M}(\gamma)$. By strict convexity of J and the assumption that $J(\bar{\nu}) = J(\bar{\nu}^\sigma)$,

$$J\left((1/2)\bar{\nu} + (1/2)\bar{\nu}^\sigma\right) < \frac{1}{2}J(\bar{\nu}) + \frac{1}{2}J(\bar{\nu}^\sigma) = J(\nu),$$

which contradicts $\bar{\nu}$ being the unique minimizer of J over $\mathcal{M}(\gamma)$. Hence, the minimizing ν must be symmetric and therefore satisfy (3.5).

It remains to show that for any $\nu \in \mathcal{M}(\gamma)$, and any permutation σ, $J(\nu) = J(\nu^\sigma)$. Let $\eta^\sigma(\mathbf{x}) = \nu^\sigma(\mathbf{x})/\bar{\mu}(\mathbf{x})$. Then $J(\nu) = J(\nu^\sigma)$ follows directly from Lemma 2.1 (symmetry properties for q^∞ and Γ^∞) together with the definition of ν^σ and the symmetry of $\bar{\mu}$:

$$\begin{aligned} J(\nu^\sigma) &= \sum_{\mathbf{x}\in\mathcal{S}^K} q^\infty(\mathbf{x})\eta^\sigma(\mathbf{x})\bar{\mu}(\mathbf{x}) - \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^K} (\eta^\sigma)^{1/2}(\mathbf{x})(\eta^\sigma)^{1/2}(\mathbf{y})\Gamma^\infty_{\mathbf{x},\mathbf{y}}\bar{\mu}(\mathbf{x}) \\ &= \sum_{\mathbf{x}\in\mathcal{S}^K} q^\infty(\mathbf{x}^\sigma)\eta(\mathbf{x}^\sigma)\bar{\mu}(\mathbf{x}^\sigma) - \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^K} \eta^{1/2}(\mathbf{x}^\sigma)\eta^{1/2}(\mathbf{y}^\sigma)\Gamma^\infty_{\mathbf{x}^\sigma,\mathbf{y}^\sigma}\bar{\mu}(\mathbf{x}^\sigma) \\ &= J(\nu). \end{aligned}$$

This completes the proof. ■

Remark 3.4 Up to this point no results have relied heavily upon the assumption of a finite state space S. Indeed, the large deviation principle of Proposition 3.1 holds in greater generality than what is presented here. For example, one can consider a compact (this condition can also be weakened) Polish space S and let jump intensities on this space to describe the dynamics of the processes. Then, in the case $K = 2$, as long as the measures μ_1 and μ_2 of interest are the unique invariant measures for the chosen jump intensities, the large deviation principle will hold. Sufficient conditions for this can be found in [5]. Moreover, under reversibility assumptions the symmetry properties proved in Lemma 2.1 also hold for the associated transition kernels and thus Proposition 3.2 remains valid as well.

4 An ergodic control problem arising from infinite swapping

The aim of this section is to introduce a finite time stochastic control problem, as well as the corresponding ergodic control problem, that are related to the infinite swapping process $\mathbf{Y}^\infty$. The control problem introduced here will

be used in the following sections to show asymptotic results regarding the particle-temperature associations and discuss the behavior and performance of infinite swapping when the underlying potential landscape exhibits asymmetry. Although all results stated here are standard in stochastic control, we include them since they may not be familiar to readers with experience in Monte Carlo methods.

With Γ^∞ the rate matrix of the infinite swapping process $\mathbf{Y}^\infty$, we simplify notation slightly by letting $r(\boldsymbol{y},\boldsymbol{z}) = \Gamma^\infty_{(\boldsymbol{y},\boldsymbol{z})}$ for $\boldsymbol{y},\boldsymbol{z} \in \mathcal{S}^2$. The generator of the process is then given by

$$\mathcal{L}f(\boldsymbol{y}) = \sum_{\boldsymbol{z}\in\mathcal{S}^2} [f(\boldsymbol{z}) - f(\boldsymbol{y})] r(\boldsymbol{y},\boldsymbol{z}),\ \boldsymbol{y} \in \mathcal{S}^2.$$

The infinite swapping process $\mathbf{Y}^\infty$ takes values in $\mathcal{D}([0,\infty) : \mathcal{S}^2)$, but for each fixed $T < \infty$ we can also consider it as an element of $\mathcal{D}([0,T] : \mathcal{S}^2)$. Our interest is now in evaluating the normalized expectation $\frac{1}{T}\log E[e^{-TF(\mathbf{Y}^\infty(\cdot))}]$, in the limit as $T \to \infty$, for functionals $F : \mathcal{D}([0,T] : \mathcal{S}^2) \to \mathbb{R}$ that are of the form

$$F(\mathbf{Y}^\infty) = \frac{1}{T}\int_0^T h(\mathbf{Y}^\infty(s))ds, \tag{4.1}$$

for some function $h : \mathcal{S}^2 \to \mathbb{R}$. The reason for studying the quantity $\frac{1}{T}\log E[e^{-TF(\mathbf{Y}^\infty(\cdot))}]$ is that, using the large deviation results of Section 3, the limit ($T \to \infty$) can be related to certain optimization problems, which in turn are of interest for evaluating the performance of infinite swapping. This is carried out in Sections 5 and 6. Throughout the section the infinite swapping process is assumed to start in some state $\mathbf{y}_0 \in \mathcal{S}^2$.

4.1 A stochastic control problem

Take $T < \infty$ to be fixed. Because the state space is finite, in formulating a stochastic control representation for $E[e^{-TF(\mathbf{Y}^\infty(\cdot))}]$ we will be able to restrict to feedback controls. The control space will be a collection of rates, and therefore takes the form $U \doteq [0,\infty)^{|\mathcal{S}|^2}$. Let $\mathcal{U}^T$ be the space of functions $u : [0,T]\times\mathcal{S}^2 \to U$ that are continuous in t and for which the $\mathbf{z}$th component of the vector $u(t,\mathbf{y})$ is positive only if $r(\mathbf{y},\mathbf{z}) > 0$:

$$\begin{aligned}\mathcal{U}^T = \big\{ u : [0,T]\times\mathcal{S}^2 \to U :\ & u(\mathbf{y},t) \text{ continuous in } t,\\ & u(t,\mathbf{y};\mathbf{z}) > 0 \text{ only if } r(\mathbf{y},\mathbf{z}) > 0, t \in [0,T]\big\},\end{aligned}$$

where $u(t,\mathbf{y};\mathbf{z})$ denotes the component of the vector $u(\mathbf{y},t) \in U$ corresponding to $\mathbf{z} \in \mathcal{S}^2$. Then to each control $u \in \mathcal{U}^T$ we will associate a controlled

process $\bar{\mathbf{Y}}^\infty$, where for each $t \in [0,T]$ and $\mathbf{y} \in \mathcal{S}^2$, the set of jump intensities for $\bar{\mathbf{Y}}^\infty$ when in state $\mathbf{y}$ at time t is given by $u(t,\mathbf{y}) \in U$; the jump intensity from $\mathbf{y}$ to $\mathbf{z} \in \mathcal{S}^2$ is $u(t,\mathbf{y};\mathbf{z})$. Although $\bar{\mathbf{Y}}^\infty$ depends on u, this is not made explicit in the notation, though the overbar indicates we consider a controlled process rather than the original infinite swapping process $\mathbf{Y}^\infty$.

We make a slight abuse of notation and also denote by u the elements in U that the control processes can take on. With this notation the generator of the controlled process $\bar{\mathbf{Y}}^\infty$ is given by $(\mathcal{L}^{u(t,\mathbf{y})}f)(\mathbf{y},t)$, where

$$(\mathcal{L}^u f)(\mathbf{y},t) = \sum_{\mathbf{z}\in\mathcal{S}^2} [f(\mathbf{z}) - f(\mathbf{y})]u(\mathbf{z}).$$

To discuss existence and uniqueness (in law) of controlled processes we use a martingale problem characterization. For a specific feedback control $u \in \mathcal{U}^T$ we say that u has an associated controlled process starting at $\mathbf{y}_0$ at time t if the following holds. On some probability space $(\Omega, \mathcal{F}, P)$, equipped with a filtration $\{\mathcal{F}_t\}_{t\in[0,T]}$, there exists a Markov process $\{\bar{\mathbf{Y}}^\infty(t) : t \in [0,T]\}$, satisfying $\bar{\mathbf{Y}}^\infty(0) = \mathbf{y}_0$ and, for $t > 0$,

$$f(t,\bar{\mathbf{Y}}^\infty(t)) - f(0,\mathbf{y}_0) - \int_0^t \left((\mathcal{L}^{u(r,\bar{\mathbf{Y}}^\infty(r))}f)(r,\bar{\mathbf{Y}}^\infty(r)) + f_t(r,\bar{\mathbf{Y}}^\infty(r))\right) dr$$

is an $\mathcal{F}_t$-martingale for all $f : [0,T] \times \mathcal{S}^2 \to \mathbb{R}$ that are bounded and continuously differentiable in t. Since $\mathcal{S}^2$ is a finite space and we consider controls that are bounded in the time variable the existence of a solution to the martingale problem is guaranteed. Indeed, for a control $u \in \mathcal{U}^T$, one can explicitly construct an associated controlled process $\bar{\mathbf{Y}}^\infty$ that solves the martingale problem by taking the correct exponential clocks etc., see [8, Chapter 4]. This process is in fact simply the jump Markov process with the given (smooth in t) jump rates. The constructed process is a solution to the associated martingale problem, and when combined with the Feller property we have that this is indeed the unique solution [14, 8].

A key ingredient in what will follow in this and the next two sections is the following stochastic control representation. For $F : \mathcal{D}([0,T] : \mathcal{S}^2) \to \mathbb{R}$ of the form (4.1),

$$\begin{aligned} &-\log E\left[e^{-TF(\mathbf{Y}^\infty)}\right] \qquad (4.2)\\ &= \inf_{u\in\mathcal{U}^T} E\left[\int_0^T \left(\sum_{\mathbf{z}\in\mathcal{S}^2 : r(\boldsymbol{z},\bar{\mathbf{Y}}^\infty(s))>0} r(\bar{\mathbf{Y}}^\infty(s),\mathbf{z})\ell\left(\frac{u(s,\bar{\mathbf{Y}}^\infty(s);\mathbf{z})}{r(\bar{\mathbf{Y}}^\infty(s),\mathbf{z})}\right) + h(\bar{\mathbf{Y}}^\infty(s))\right) ds\right], \end{aligned}$$

where $\bar{\mathbf{Y}}^\infty(0) = \mathbf{y}_0$, the infimum is over all controls u, and ℓ is the function

$$\ell(x) = \begin{cases} x \log x - x + 1, & x \geq 0, \\ \infty, & \text{otherwise.} \end{cases}$$

The right-hand side of (4.2) is a stochastic control problem with running cost $c : \mathcal{S}^2 \times U \to \mathbb{R}$ given by

$$c(\mathbf{y}, u) = \sum_{\mathbf{z} \in \mathcal{S}^2 : r(\mathbf{y},\mathbf{z}) > 0} r(\mathbf{y}, \mathbf{z}) \ell \left(\frac{u(\mathbf{z})}{r(\mathbf{y}, \mathbf{z})} \right) + h(\mathbf{y}). \tag{4.3}$$

Representations such as (4.2) are commonly used in connection with large deviations and similar results can be found in, e.g., [**?**, **?**, **?**, 9].

To discuss the dynamic programming equation associated with the stochastic control problem in (4.2), define $\bar{W}^T(t, \mathbf{y})$ to be the conditional version of the right-hand side of the representation:

$$\bar{W}^T(t, \mathbf{y}) \doteq \inf_{u \in \mathcal{U}^T} E_{t,\mathbf{y}} \left[\int_t^T c \left(\bar{\mathbf{Y}}^\infty(s), u(s, \bar{\mathbf{Y}}^\infty(s)) \right) \right], \tag{4.4}$$

where $E_{t,\mathbf{y}}$ denotes conditional expectation with respect to $\bar{\mathbf{Y}}^\infty(t) = \mathbf{y}$. Note that the representation (4.2) then involves the conditional expectation $E_{0,\mathbf{y}_0}$. Proposition 4.1, which follows from a standard verification argument, shows that any C^1 solution of the dynamic programming equation

$$W_t^T(t, \mathbf{y}) + \inf_{u \in U} \left\{ \mathcal{L}^u W^T(t, \mathbf{y}) + c(\mathbf{y}, u) \right\} = 0, \tag{4.5}$$

with the associated terminal condition $W^T(T, \mathbf{y}) = 0$ for $\mathbf{y} \in \mathcal{S}^2$, is equal to $\bar{W}^T$.

Before stating and proving the verification theorem just eluded to, we consider the optimal control in (4.5). From the definition of $\mathcal{L}^u$ it follows that, for any $u \in U$,

$$\mathcal{L}^u W^T(t, \mathbf{y}) = \sum_{\mathbf{z} \in \mathcal{S}^2} \left[W^T(t, \mathbf{z}) - W^T(t, \mathbf{y}) \right] u(\mathbf{z}).$$

Straightforward calculus shows that the minimizing vector u in (4.5) is given by

$$\bar{u}(t, \mathbf{y}; \mathbf{z}) = r(\mathbf{y}, \mathbf{z}) e^{-(W^T(t,\mathbf{z}) - W^T(t,\mathbf{y}))}, \ \mathbf{z} \in \mathcal{S}^2. \tag{4.6}$$

Inserting this expression into (4.5) gives a total of N^2 coupled differential equations for W^T, one for each $\mathbf{y} \in \mathcal{S}^2$, which we write as

$$W_t^T(t, \mathbf{y}) = -\sum_{z \in \mathcal{S}^2} r(\mathbf{y}, \mathbf{z}) \left(1 - e^{-(W^T(t,\mathbf{z}) - W^T(t,\mathbf{y}))} \right) - h(\mathbf{y}). \tag{4.7}$$

We are now ready to state the relevant verification result. The argument is standard, but is included for completeness.

Proposition 4.1 *Suppose that $W^T : [0,T] \times \mathcal{S}^2 \to \mathbb{R}$ is a C^1 (in t) solution to the dynamic programming equation (4.5). Define $\bar{u}(t,\mathbf{y};\mathbf{z}) \in \mathcal{U}^T$ by (4.6), and let $\bar{\mathbf{Y}}^\infty$ be the corresponding jump Markov process (i.e., solution to the martingale problem). Then W^T equals the minimal cost function $\bar{W}^T$ defined in (4.4), and*

$$W^T(t,\mathbf{y}) = E_{t,\mathbf{y}}\left[\int_t^T c\left(\bar{\mathbf{Y}}^\infty(s), \bar{u}(s,\bar{\mathbf{Y}}^\infty(s))\right) ds\right],\ (t,\mathbf{y}) \in [0,T]\times\mathcal{S}^2,$$

so that $\bar{u} \in \mathcal{U}^T$ is the optimal control and $\bar{\mathbf{Y}}^\infty$ is the optimally controlled process.

Proof. To emphasize the choice of control, let E^u denote expectation when the control $u \in \mathcal{U}^T$ is used and take $\tilde{u} \in \mathcal{U}^T$ to be any control with associated controlled process $\tilde{\mathbf{Y}}^\infty$. From the martingale property it follows that

$$\begin{aligned}
&E^{\tilde{u}}_{t,\mathbf{y}}\left[W^T(T,\tilde{\mathbf{Y}}^\infty(T))\right] \\
&\quad = W^T(t,\mathbf{y}) + E^{\tilde{u}}_{t,\mathbf{y}}\left[\int_t^T \left(\mathcal{L}^{\tilde{u}(s,\tilde{\mathbf{Y}}^\infty(s))} W^T(s,\tilde{\mathbf{Y}}^\infty(s)) + W^T_t(s,\tilde{\mathbf{Y}}^\infty(s))\right) ds\right].
\end{aligned}$$

The terminal condition is $W^T(T,\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathcal{S}^2$ and the left-hand side is thus 0. Moreover, since W^T solves the dynamic programming equation (4.5) we have

$$W^T_t(t,\mathbf{y}) + \mathcal{L}^{\tilde{u}(t,\mathbf{y})} W^T(t,\mathbf{y}) \geq -c(\mathbf{y},\tilde{u}(t,\mathbf{y})),$$

and therefore obtain a lower bound on the cost:

$$W^T(t,\mathbf{y}) \leq E^{\tilde{u}}_{t,\mathbf{y}}\left[\int_t^T c(\tilde{\mathbf{Y}}^\infty(s), \tilde{u}(s,\tilde{\mathbf{Y}}^\infty(s)))ds\right].$$

If we instead use the optimal control $\bar{u}$, for which the infimum in (4.5) is attained, all inequalities become equalities, and therefore

$$W^T(t,\mathbf{y}) = E^{\bar{u}}_{t,\mathbf{y}}\left[\int_t^T c(\bar{\mathbf{Y}}^\infty(s), \bar{u}(s,\bar{\mathbf{Y}}^\infty(s)))ds\right].$$

This shows that W^T is indeed the minimal cost function and $\bar{u}$ is the optimal control. ■

Before taking the limit as $T \to \infty$, suppose we have a solution W^T to (4.5) and consider the function $V : [0,T] \times \mathcal{S}^2 \to \mathbb{R}$ defined by $V(t,\mathbf{y}) \doteq e^{-W^T(t,\mathbf{y})}$. Then W^T satisfying (4.7) implies that V satisfies the system

$$0 = V_t(t,\mathbf{y}) + \sum_{\mathbf{z}\in\mathcal{S}^2} r(\mathbf{y},\mathbf{z})\left(V(t,\mathbf{z}) - V(t,\mathbf{y})\right) - h(\mathbf{y})V(t,\mathbf{y}),$$

with terminal condition $V(T,\mathbf{y}) = 1$. This is a finite system of linear ordinary differential equations and existence and uniqueness of a solution V, and thus W^T, hold. Moreover, it can be shown that in fact $V(t,\mathbf{y}) = E_{t,\mathbf{y}}[\exp - \int_t^T h(\mathbf{Y}^\infty(s))ds]$ and, since $V(0,0)$ and $W^T(0,0)$ correspond to the two quantities in (4.2), uniqueness of the solution of the system of ODEs together with Proposition 4.1 then confirm the representation (4.2).

4.2 Limit control problem as $T \to \infty$

As mentioned at the beginning of this section, we are ultimately interested in the limit as $T \to \infty$ of the normalized expectation $-\frac{1}{T}\log E[e^{-TF(\mathbf{Y}^\infty)}]$. Given the representation (4.2) this is equivalent to taking the limit of

$$\inf_{u\in\mathcal{U}^T} E\left[\frac{1}{T}\int_0^T \left(\sum_{z\in\mathcal{S}^2} r(\bar{\mathbf{Y}}^\infty(s),z)\ell\left(\frac{u(s,\bar{\mathbf{Y}}^\infty(s);z)}{r(\bar{\mathbf{Y}}^\infty(s),z)}\right) + h(\bar{\mathbf{Y}}^\infty(s))\right) ds\right],$$

where $\bar{\mathbf{Y}}^\infty$ is the controlled process associated with control u and the infimum is over all such feedback controls. In light of the previous subsection, this is precisely $\lim_{T\to\infty} W^T(0,\mathbf{y}_0)/T$, for $\mathbf{y}_0 \in \mathcal{S}^2$, which falls under the umbrella of ergodic control problems, or "average cost per unit time"; some general references are [10, 12]. Since the set of initial conditions is finite convergence will be uniform with respect to this parameter, and hence it is not made explicit in the notation.

The limit Bellman equation is

$$\inf_{u\in U}\{\mathcal{L}^u W(y) - \gamma + c(\mathbf{y},u)\} = 0, \tag{4.8}$$

where W and γ are unknown, with γ the sought-after limit, and c is defined in (4.3). As is well known the solution W to such an equation is unique only up to an additive constant. Together with the form of the generator $\mathcal{L}^u$, the definition of c implies that, for each $\mathbf{y} \in \mathcal{S}^2$, the Bellman equation (4.8)

takes the form

$$\inf_{u\in U}\left\{\sum_{\mathbf{z}\in\mathcal{S}^2}\left(u(\mathbf{z})[W(\mathbf{z})-W(\mathbf{y})]+r(\mathbf{y},\mathbf{z})\ell\left(\frac{u(\mathbf{z})}{r(\mathbf{y},\mathbf{z})}\right)\right)-\gamma+h(\mathbf{y})\right\}=0. \tag{4.9}$$

The minimizing $\bar{u}$ takes the same form as for the pre-limit problem,

$$\bar{u}(\mathbf{z})=r(\mathbf{y},\mathbf{z})e^{-(W(\mathbf{z})-W(\mathbf{y}))}, \tag{4.10}$$

with the corresponding generator

$$\mathcal{L}^{\bar{u}(\mathbf{z})}f(y)=\sum_{\mathbf{z}\in\mathcal{S}^2}r(\mathbf{y},\mathbf{z})e^{-(W(\mathbf{z})-W(\mathbf{y}))}[f(\mathbf{z})-f(\mathbf{y})].$$

Similar to the equation for the pre-limit control problem in the previous subsection, inserting the optimal $\bar{u}$ in (4.8) yields the equation

$$0=\sum_{z\in\mathcal{S}^2}r(\mathbf{y},\mathbf{z})[1-e^{-(W(\mathbf{z})-W(\mathbf{y}))}]-\gamma+h(\mathbf{y}). \tag{4.11}$$

The main facts we will need regarding this problem are the following. Under our conditions, which include the ergodicity of the dynamics in the original infinite swapping process $\mathbf{Y}^\infty$, a solution (γ, W) to the dynamic programming equation (4.11) exists and is unique (up to an additive constant in W), and (4.7) defines and optimal control. The proof of the first statement follows from classical arguments based on approximation by so-called "discounted" control problems, and the second follows from a verification argument very much like the one used in the last section for the corresponding finite time problem.

Using the reversibility of the infinite swapping dynamics (discussed in Section 3) shows that the invariant measure $\bar{\nu}$ for the optimally controlled process is

$$\bar{\nu}(x_1,x_2)=\bar{\mu}(x_1,x_2)e^{-2W(x_1,x_2)-a},\ (x_1,x_2)\in\mathcal{S}^2,$$

where a is a normalizing constant. However, W is unique only up to an additive constant, and so we can assume without loss that for $a=0$, $\bar{\nu}$ defines a probability measure on $\mathcal{S}^2$.

5 A diagnostic for the convergence of the empirical measure

One of the challenges of Monte Carlo when dealing with problems involving rare events is to determine when the algorithm has converged. For example,

in the setting of MCMC it can happen that the empirical measure appears to have converged, when in reality the underlying process is stuck in some collection of metastable states, and significant parts of the state space have not been visited nearly often enough for a good approximation to the true equilibrium distribution. Hence it is of interest to know if there are diagnostics that can determine when convergence has or has not taken place.

In this section we rigorously justify a diagnostic that will tell the user when INS has *not* converged. More precisely, we will show using a large deviations analysis that the empirical measure of the particle/temperature association introduced in Section 3 provides such a diagnostic, as do various functionals of this empirical measure (see Remark 5.2). In particular, ρ_T must converge to the uniform distribution on Σ_K if the weighted empirical measure η_T^∞ is to converge to the true stationary distribution. Although the convergence of η_T^∞ is what is needed for computational purposes, the functionals of the empirical measure ρ_T can be readily observed while running a simulation, and therefore provide convenient diagnostics.

The main result of this section, Proposition 5.1, gives the precise relation between the convergences of ρ_T and η_T^∞. The result is of an asymptotic character and indicates how convergence of the empirical measure η_T^∞ to μ can only occur if there is also convergence of ρ_T to the uniform distribution. Since $\mathcal{P}(\Sigma_2)$ is a finite dimensional space, convergence in the weak topology is the same as ordinary convergence as elements of a subset of a Euclidean space. We also consider $\mathcal{P}(\mathcal{S}^2)$ with a metric that is consistent with weak convergence and under which it is a Polish space. In the statement of the proposition $\mathcal{N}_a(w^*)$ denotes the open neighborhood about w^* of radius a in $\mathcal{P}(\Sigma_2)$, and similarly for $\mathcal{N}_\epsilon(\mu) \subset \mathcal{P}(\mathcal{S}^2)$.

Proposition 5.1 *Let $w^* = (1/2, 1/2)$. Then for each $a > 0$ there is an $\epsilon > 0$ such that*

$$P\left(\eta_T^\infty \in \mathcal{N}_\epsilon(\mu) | \rho_T \in (\mathcal{N}_a(w^*))^c\right) \to 0 \text{ as } T \to \infty.$$

Remark 5.2 Although the proof of Proposition 5.1 is given for the case of two temperatures, the analogous result for any finite number of temperatures holds, though the notation needed for the proof is more complicated. In this more general setting, it is worth noting that any functional of ρ_T must also converge to its asymptotic counterpart before η_T^∞ can converge to μ. Thus one can consider diagnostics that are based on lower dimensional quantities. For example, in place of the $K!$-dimensional object $\rho_T = \{\rho_T^\sigma\}_{\sigma \in \Sigma_K}$, one

could use the K-dimensional object defined by

$$[\beta_T]_k \doteq \sum_{\sigma\in\Sigma_K:\sigma(1)=k} \rho_T^\sigma, \quad k=1,\dots,K.$$

This quantity can be interpreted as follows, using the convention for particle and temperature associations when using INS that was introduced previously. Let each particle be associated with the temperature (dynamic) it was assigned at time zero. Then $[\beta_T]_k$ is the fraction of time the particle initially assigned temperature 1 uses temperature k in $[0,T]$. Since $\rho_T^\sigma \to 1/K!$ and particles are exchangeable, $[\beta_T]_k \to 1/K$ as $T\to\infty$. It is in fact this diagnostic that has been used in previous numerical studies such as [2, 1].

To prove Proposition 5.1 we first study the probability measure that minimizes the large deviation rate I for any fixed $\bar{w}\in\mathcal{P}(\Sigma_K)$, with $\bar{w}$ not the uniform distribution.

Lemma 5.3 *For any $\bar{w}\in\mathcal{P}(\Sigma_2)$ not equal to the uniform distribution, the infimum of $I(\gamma,\bar{w})$ over γ is uniquely attained at some $\bar{\gamma}\neq\mu$.*

The proof will use the ergodic control problem of the previous section. In this section we give the proof for $K=2$, and outline the proof for general K in Remark A.4 in the appendix. The following lemma ensures that we can switch our focus from the unconstrained version of the optimization problem of Lemma 5.3 to a related ergodic control problem, and that there is a correspondence between the minimizers in the two settings. Recall the solution (γ, W) to the Bellman equation (4.9) has the property that γ is unique and W is unique up to an additive constant (see, e.g., [12, Chapter 7]).

Lemma 5.4 *Consider the optimization problem*

$$\inf_{\nu\in\mathcal{P}(\mathcal{S}^2)} \left\{ J(\nu) + \sum_{\mathbf{x}\in\mathcal{S}^2} h(\mathbf{x})\nu(\mathbf{x}) \right\}, \tag{5.1}$$

and also the static Bellman equation given in (4.9) *in Section 4:*

$$0 = \inf_{u\in U} \left\{ \sum_{\mathbf{y}\in\mathcal{S}^2} \left(u(\mathbf{y})\left[W(\mathbf{y})-W(\mathbf{x})\right] + r(\mathbf{x},\mathbf{y})\ell\left(\frac{u(\mathbf{y})}{r(\mathbf{x},\mathbf{y})}\right)\right) - \gamma + h(\mathbf{x}) \right\}.$$

Consider any solution (γ, W) to the Bellman equation, with W taken to be normalized in the sense that

$$\bar{\nu}(\mathbf{x}) = \bar{\mu}(\mathbf{x})e^{-2W(\mathbf{x})} \tag{5.2}$$

is a probability measure on $\mathbf{S}^2$. Then $\bar{\nu}$ is a minimizer in (5.1). Conversely, consider a minimizing measure ν^* in (5.1). Then a solution (γ^*, W^*) to the Bellman equation is given by

$$W^*(\mathbf{x}) = -\log\left[\frac{d\nu^*}{d\bar{\mu}}\right]^{1/2}(\mathbf{x}),\ \gamma^* = J(\nu^*) + \sum_{\mathbf{x}\in\mathcal{S}^2} h(\mathbf{x})\nu^*(\mathbf{x}),$$

and by uniqueness $(\gamma^*, W^*) = (\gamma, W)$.

Proof. We start by showing that an averaged version of the static Bellman equation gives an upper bound for the optimization problem. Suppose (γ, W) is a solution to the Bellman equation and define the measure $\bar{\nu}$ by (5.2). The Bellman equation holds with equality for all $\mathbf{x}$, and averaging with respect to $\bar{\nu}$ gives

$$\gamma = \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^2} r(\mathbf{x},\mathbf{y})\left[1 - e^{-(W(\mathbf{y})-W(\mathbf{x}))}\right]\bar{\mu}(\mathbf{x})e^{-2W(\mathbf{x})} + \sum_{\mathbf{x}\in\mathcal{S}^2} h(\mathbf{x})\bar{\mu}(\mathbf{x})e^{-2W(\mathbf{x})}. \tag{5.3}$$

Define $\theta(\mathbf{x}) = e^{-2W(\mathbf{x})}$, the likelihood ratio of $\bar{\nu}$ and $\bar{\mu}$, and consider the rate function J evaluated at $\bar{\nu}$:

$$\begin{aligned} J(\bar{\nu}) &= \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^2}\left(\theta(\mathbf{x}) - \theta^{1/2}(\mathbf{x})\theta^{1/2}(\mathbf{y})\right) r(\mathbf{x},\mathbf{y})\bar{\mu}(\mathbf{x}) \\ &= \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^2}\left(e^{-2W(\mathbf{x})} - e^{-W(\mathbf{x})-W(\mathbf{y})}\right) r(\mathbf{x},\mathbf{y})\bar{\mu}(\mathbf{x}) \\ &= \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^2} r(\mathbf{x},\mathbf{y})\left(1 - e^{-(W(\mathbf{y})-W(\mathbf{x}))}\right)\bar{\mu}(\mathbf{x})e^{-2W(\mathbf{x})}. \end{aligned}$$

Then using (5.3) for the last equality,

$$\begin{aligned} &\inf_{\nu\in\mathcal{P}(\mathcal{S}^2)}\left\{J(\nu) + \sum_{\mathbf{x}\in\mathcal{S}^2} h(\mathbf{x})\nu(\mathbf{x})\right\} \\ &\quad\le \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^2} r(\mathbf{x},\mathbf{y})\left(1 - e^{-(W(\mathbf{y})-W(\mathbf{x}))}\right)\bar{\mu}(\mathbf{x})e^{-2W(\mathbf{x})} + \sum_{\mathbf{x}\in\mathcal{S}^2} h(\mathbf{x})e^{-2W(\mathbf{x})}\bar{\mu}(\mathbf{x}) \\ &\quad= \gamma. \end{aligned}$$

Next define γ^* to be the minimal value

$$\gamma^* = \inf_{\nu \in \mathcal{P}(\mathcal{S}^2)} \left\{ J(\nu) + \sum_{\mathbf{x} \in \mathcal{S}^2} h(\mathbf{x}) \nu(\mathbf{x}) \right\}.$$

From the previous display $\gamma^* \leq \gamma$, and we now proceed to show the reverse inequality.

Let ν^* denote a minimizing measure, i.e.,

$$\nu^* = \arg\min_{\nu \in \mathcal{P}(\mathcal{S}^2)} \left\{ J(\nu) + \sum_{\mathbf{x} \in \mathcal{S}^2} h(\mathbf{x}) \nu(\mathbf{x}) \right\}.$$

The existence of such a measure follows from the fact that J has compact level sets and the boundedness of h. Moreover strict convexity of J implies it is unique. Define $\theta^* = [d\nu^*/d\bar{\mu}]$. For any $\mathbf{x} \in \mathcal{S}^2$ we have $\theta^*(\mathbf{x}) \in (0, \infty)$. For the upper bound, note that

$$\theta^*(\mathbf{x}) \leq \frac{\max_{\mathbf{x} \in \mathcal{S}^2} \nu^*(\mathbf{x})}{\min_{\mathbf{x} \in \mathcal{S}^2} \bar{\mu}(\mathbf{x})} < \infty.$$

The second inequality is due to ν^* being a probability measure and the fact that $\bar{\mu}$ has support $\mathcal{S}^2$, which implies that $\bar{\mu}(\mathbf{x}) > 0$ for all $\mathbf{x} \in \mathcal{S}^2$. It follows from the finiteness of $\mathcal{S}^2$ that θ^* is bounded from above. Moreover, by differentiating the objective function it is not difficult to check that the optimal choice θ^* will satisfy

$$\theta^*(\mathbf{x}) = \frac{1}{4(\sum_{\mathbf{y} \in \mathcal{S}^2} r(\mathbf{x}, \mathbf{y}) + h(\mathbf{x}))^2} \left(\sum_{\mathbf{y} \in \mathcal{S}^2} r(\mathbf{x}, \mathbf{y}) \left(\theta^*\right)^{1/2} (\mathbf{y}) \right)^2, \ \mathbf{x} \in \mathcal{S}^2.$$

Suppose that $\theta^*(\mathbf{x})$ is zero for at least one $\mathbf{x} \in \mathcal{S}^2$. Since $\sum_{\mathbf{y} \in \mathcal{S}^2} r(\mathbf{x}, \mathbf{y}) + h(\mathbf{x}) \in (0, \infty)$ for all $\mathbf{x}$ for $\theta^*(\mathbf{x})$ to be zero it must hold that the sum in the last display is zero. The underlying jump rates are such that $\mathcal{S}^2$ forms a communicating class and thus for each $\mathbf{x}$ there is at least one $\mathbf{y}$ such that $r(\mathbf{x}, \mathbf{y}) > 0$. It follows that $\theta^*(\mathbf{x}) = 0$ requires $\theta^*(\mathbf{y}) = 0$ for all $\mathbf{y}$ with which $\mathbf{x}$ communicates. Repeating this argument, using that $\mathcal{S}^2$ is a communicating class under the original dynamics, shows that if $\theta^*(\mathbf{x}) = 0$, then $\theta^* \equiv 0$. This is clearly a contradiction and it must hold that $\theta^*(\mathbf{x}) > 0$ for all $\mathbf{x}$. Hence, $\theta^*(\mathbf{x}) \in (0, \infty)$ for all $\mathbf{x}$.

Set

$$W^*(\mathbf{x}) = -\log(\theta^*)^{1/2}(\mathbf{x}).$$

Inserting the measure ν^* into the objective function and rewriting it in terms of W^*,

$$\begin{aligned}\gamma^* &= J(\nu^*) + \sum_{\mathbf{x}\in\mathcal{S}^2} h(\mathbf{x})\nu^*(\mathbf{x}) \\ &= \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^2} \left(e^{-2W^*(\mathbf{x})} - e^{-W^*(\mathbf{y})}e^{-W^*(\mathbf{x})}\right) r(\mathbf{x},\mathbf{y})\bar{\mu}(\mathbf{x}) + \sum_{\mathbf{x}\in\mathcal{S}^2} h(\mathbf{x})\theta^*(\mathbf{x})\bar{\mu}(\mathbf{x}) \\ &= \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^2} r(\mathbf{x},\mathbf{y})\left(1 - e^{-(W^*(\mathbf{y})-W^*(\mathbf{x}))}\right)\bar{\mu}(\mathbf{x})e^{-2W^*(\mathbf{x})} + \sum_{\mathbf{x}\in\mathcal{S}^2} h(\mathbf{x})\bar{\mu}(\mathbf{x})e^{-2W^*(\mathbf{x})}.\end{aligned}$$

This is the Bellman equation averaged with respect to $\bar{\mu}(\mathbf{x})e^{-2W^*(\mathbf{x})}$.

The cost (average cost per unit time) associated with the control u^* is (see Section 4)

$$\lim_{T\to\infty} E\left[\frac{1}{T}\int_0^T \left(\sum_{\mathbf{y}\in\mathcal{S}^2} r(\bar{\mathbf{Y}}^\infty(s),\mathbf{y})\ell\left(\frac{u^*(s,\bar{\mathbf{Y}}^\infty(s);\mathbf{y})}{r(\bar{\mathbf{Y}}^\infty(s),\mathbf{y})}\right) + h(\bar{\mathbf{Y}}^\infty(s))\right) ds\right].$$

where the controlled process $\bar{\mathbf{Y}}^\infty$ has dynamics according to the choice of control u^*. Using the same calculations as in Section 4, the invariant measure associated with this process is precisely ν^* and by ergodicity the limit in the last display is the average of the cost with respect to ν^*:

$$\begin{aligned}&\sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^2} \left(r(\mathbf{x},\mathbf{y})\ell\left(\frac{u^*(\mathbf{y})}{r(\mathbf{x},\mathbf{y})}\right) + h(\mathbf{x})\right)\nu^*(\mathbf{x}) \\ &\quad= \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^2} \left(r(\mathbf{x},\mathbf{y})\left[1 - e^{-W^*(\mathbf{y})+W^*(\mathbf{x})}\right] + h(\mathbf{x}) - r(\mathbf{x},\mathbf{y})\left(W^*(\mathbf{y}) - W^*(\mathbf{x})\right)\right)\nu^*(\mathbf{x}) \\ &\quad= \sum_{\mathbf{x},\mathbf{y}\in\mathcal{S}^2} \left(r(\mathbf{x},\mathbf{y})\left[1 - e^{-W^*(\mathbf{y})+W^*(\mathbf{x})}\right] + h(\mathbf{x})\right)\nu^*(\mathbf{x}) - \sum_{\mathbf{x}\in\mathcal{S}^2} \mathcal{L}^{u^*}W^*(\mathbf{x})\nu^*(\mathbf{x}).\end{aligned}$$

The first term is γ^*. Moreover, since $\mathcal{L}^{u^*}$ is the generator associated with ν^*, the second term is 0 by Echeverria's theorem [8]. Thus, γ^* is also the cost obtained using the control u^*. Since γ is the optimal cost it follows that $\gamma^* \geq \gamma$.

Combining the two inequalities gives $\gamma = \gamma^*$. This implies that the two measures $\bar{\nu}$ and ν^* are both minimizers in (5.1). Strict convexity of the rate function J then ensures that the two measures are in fact the same, $\bar{\nu} = \nu^*$ and by extension $W = W^*$. ■

Proof of Lemma 5.3. Fix a $\bar{w} \in \mathcal{P}(\Sigma_2) \setminus \{(1/2, 1/2)\}$ and consider the

optimization problem

$$\inf\{I(\gamma,\bar{w}):\ \gamma\in\mathcal{P}(\mathcal{S}^2)\} \tag{5.4}$$
$$=\inf_{\nu\in\mathcal{P}(\mathcal{S}^2)}\left\{J(\nu):\ \sum_{\mathbf{x}\in\mathcal{S}^2}\rho(\mathbf{x})\nu(\mathbf{x})=\bar{w}_1,\ \sum_{\mathbf{x}\in\mathcal{S}^2}\rho(\mathbf{x}^R)\nu(\mathbf{x})=\bar{w}_2\right\}.$$

We are only interested in the minimizing measure ν and it is enough to consider any optimization problem that will have the same minimizer as (5.4). Using Lagrange multipliers λ_1,λ_2, (5.4) can be formulated as the unconstrained optimization problem

$$\min_{\gamma,w}\ \{I(\gamma,w)+\lambda_1(w_1-\bar{w}_1)+\lambda_2(w_2-\bar{w}_2)\};$$

see, e.g., Theorem 8.1 and its extension to equality constraints in [13] for the existence of multipliers λ_1,λ_2. Using the fact that necessarily $\bar{w}_1+\bar{w}_2=w_1+w_2=1$, this has the same minimizer as $\min_{\gamma,w}\ \{I(\gamma,w)+(\lambda_1-\lambda_2)(w_1-\bar{w}_1)\}$, where the multipliers are chosen so that $w_1=\bar{w}_1$. With such multipliers given (and fixed), by using the definition of $I(\gamma,w)$ the optimization problem becomes

$$\min_{\nu\in\mathcal{P}(\mathcal{S}^2)}\left\{J(\nu)+(\lambda_1-\lambda_2)\sum_{\mathbf{x}\in\mathcal{S}^2}\rho(\mathbf{x})\nu(\mathbf{x})\right\}-(\lambda_1-\lambda_2)\bar{w}_1,$$

and we further simplify by dropping the term $-(\lambda_1-\lambda_2)\bar{w}_1$.

The Lagrange multipliers λ_1 and λ_2 correspond to $\bar{w}\neq(1/2,1/2)$ so it cannot be the case that $\lambda_1=\lambda_2$. Let λ denote the difference $\lambda_1-\lambda_2$; without loss of generality we can assume that $\lambda>0$. Thus, in order to prove the claim it is enough to consider the minimizer of

$$\min_{\nu\in\mathcal{P}(\mathcal{S}^2)}\left\{J(\nu)+\sum_{\mathbf{x}\in\mathcal{S}^2}h(\mathbf{x})\nu(\mathbf{x})\right\}, \tag{5.5}$$

where $h(\mathbf{x})=\lambda\rho(\mathbf{x})$.

For any cost function h we can define the functional $F:\mathcal{P}(\mathcal{S}^2)\to\mathbb{R}$ by $F(\nu)\doteq\sum_{\mathbf{x}\in\mathcal{S}^2}h(\mathbf{x})\nu(\mathbf{x})$. This choice of F is of the form considered in Section 4. In particular, F is bounded and continuous, and from the Laplace

principle for ν_T it follows that

$$\lim_{T\to\infty} -\frac{1}{T}\log E\left[e^{-TF(\nu_T)}\right] = \inf_{\nu\in\mathcal{P}(\mathcal{S}^2)}\{J(\nu)+F(\nu)\}$$
$$= \inf_{\nu\in\mathcal{P}(\mathcal{S}^2)}\left\{J(\nu)+\lambda\sum_{\mathbf{x}\in\mathcal{S}^2}\rho(\mathbf{x})\nu(\mathbf{x})\right\}.$$

In Lemma 5.4 it was shown that the minimizer of (5.5) is

$$\bar{\nu}(\mathbf{x}) = \bar{\mu}(\mathbf{x})e^{-2W(\mathbf{x})},$$

where (γ, W) is a solution to the Bellman equation

$$0 = \sum_{\mathbf{y}\in\mathcal{S}^2} r(\mathbf{x},\mathbf{y})\left[1-e^{-W(\mathbf{x})+W(\mathbf{y})}\right] - \gamma + \lambda\rho(\mathbf{x}),\ \mathbf{x}\in\mathcal{S}^2, \tag{5.6}$$

with the value function W normalized to make $\bar{\nu}$ a probability measure.

Next we use the fact that for this particular choice of $h(\mathbf{x})$, $\gamma < \lambda/2$ (see Lemma A.2 in the Appendix). The intuition here is that when $\rho \equiv 1/2$ there is no incentive to use the control and we get the cost $\gamma = \lambda/2$. Moreover when ρ depends on $\mathbf{x}$ not using an active control will result in the same cost, since $\rho(\mathbf{x}) + \rho(\mathbf{x}^R) = 1$ for all $\mathbf{x}$. In contrast, by accepting a small increase in the cost due to active control we can lower the running cost substantially by favoring states with lower ρ value than their symmetric counterpart.

The aim is to show that the optimal measure $\bar{\nu}$ is such that $M\bar{\nu} \neq \mu$. From the definitions of M and ρ [given for two temperatures in (2.5) and (3.1), respectively] for $\mathbf{x}\in\mathcal{S}^2$

$$\begin{aligned}(M\bar{\nu})(\mathbf{x}) &= \rho(\mathbf{x})\left(\bar{\nu}(\mathbf{x})+\bar{\nu}(\mathbf{x}^R)\right)\\ &= \rho(\mathbf{x})\left(\bar{\mu}(\mathbf{x})e^{-2W(\mathbf{x})}+\bar{\mu}(\mathbf{x}^R)e^{-2W(\mathbf{x}^R)}\right)\\ &= \mu(\mathbf{x})\frac{e^{-2W(\mathbf{x})}+e^{-2W(\mathbf{x}^R)}}{2}.\end{aligned}$$

We argue by contradiction. Suppose that $M\bar{\nu} = \mu$. Then by the last display

$$e^{-2W(\mathbf{x})} + e^{-2W(\mathbf{x}^R)} - 2 = 0 \text{ for } \mathbf{x}\in\mathcal{S}^2. \tag{5.7}$$

Let $\mathcal{D}$ denote the set of diagonal states in $\mathcal{S}^2$:

$$\mathcal{D} \doteq \{\mathbf{x}\in\mathcal{S}^2 :\ \mathbf{x} = \mathbf{x}^R\}.$$

Note that for $\mathbf{x} \in \mathcal{D}$ always $\rho(\mathbf{x}) = 1/2$ and hence $h(\mathbf{x}) = \lambda/2$. Moreover, as a special case of (5.7), the value function W must be zero on the diagonal $\mathcal{D}$.

The symmetrized dynamics of the INS process (see Lemma 2.1) imply that if $\mathbf{x} \in \mathcal{D}$ communicates directly with a state $\mathbf{y}$ then it communicates directly with $\mathbf{y}^R$ as well. To simplify notation we therefore let $\mathcal{A}$ denote the collection of states that lie above the diagonal. For states $\mathbf{x} \in \mathcal{D}$ the Bellman equation (5.6) then takes the form

$$0 = \sum_{\mathbf{y} \in \mathcal{A}} \left(r(\mathbf{x}, \mathbf{y}) \left[1 - e^{-(W(\mathbf{y}) - W(\mathbf{x}))} \right] + r(\mathbf{x}, \mathbf{y}^R) \left[1 - e^{-(W(\mathbf{y}^R) - W(\mathbf{x}))} \right] \right) - \gamma + \frac{\lambda}{2}.$$

By symmetry of the rates r, if $\mathbf{x} \in \mathcal{D}$ then $r(\mathbf{x}, \mathbf{y}) = r(\mathbf{x}, \mathbf{y}^R)$. Combined with the constraint that W is zero on $\mathcal{D}$ we can rewrite the Bellman equation for $\mathbf{x} \in \mathcal{D}$ as

$$0 = \sum_{\mathbf{y} \in \mathcal{A}} r(\mathbf{x}, \mathbf{y}) \left(2 - e^{-W(\mathbf{y})} - e^{-W(\mathbf{y}^R)} \right) - \gamma + \frac{\lambda}{2}.$$

The assumptions on τ_1, τ_2 and the potential V ensure that $\rho(\mathbf{x})$ is not identically equal to $1/2$, and therefore by Lemma A.2 in the Appendix $\gamma < \lambda/2$. It follows that W satisfies

$$\sum_{\mathbf{y} \in \mathcal{A}} r(\mathbf{x}, \mathbf{y}) \left(2 - e^{-W(\mathbf{y})} - e^{-W(\mathbf{y}^R)} \right) = \gamma - \frac{\lambda}{2} < 0, \ \mathbf{x} \in \mathcal{D}. \tag{5.8}$$

The rates $r(\mathbf{x}, \mathbf{y})$ are all nonnegative and for (5.8) to hold requires that for at least one of the states $\mathbf{y} \in \mathcal{A}$

$$e^{-W(\mathbf{y})} + e^{-W(\mathbf{y}^R)} > 2.$$

However, this inequality is not compatible with (5.7) [see Lemma A.1 in the Appendix with $K = 2$, $a_1 = e^{-W(\mathbf{y})}$ and $a_2 = e^{-W(\mathbf{y}^R)}$]. Hence, the condition (5.7) violates the Bellman equation for diagonal states and it cannot be that $M\bar{\nu} = \mu$. This completes the proof. ■

Proof of Proposition 5.1. In addition to showing the claimed convergence as $T \to \infty$ we will show that the probability decays exponentially in T. Consider the mapping

$$a \to R(a) \doteq \inf \left\{ I(\gamma, w) : w \in (\mathcal{N}_a(w^*)^c) \right\}.$$

The rate function J has $\bar{\mu}$, the symmetrized version of the original stationary distribution, as its unique minimizer and $J(\bar{\mu}) = 0$. Moreover, $\bar{\mu}$ maps to w^* in the sense that $\langle \rho^\sigma, \bar{\mu} \rangle = 1/2$ for $\sigma = \{1, 2\}$ and $\sigma = \{2, 1\}$ (the permutations available for two temperatures); see Section 3. Consider the set

$$C_a \doteq \{(\gamma, w) : w \in (\mathcal{N}_a(w^*))^c\}.$$

This is a closed set and, since I is a rate function on $\mathcal{P}(\mathcal{S}^2) \times \mathcal{P}(\Sigma_2)$, the infimum of I over C_a is achieved and necessarily $R(a) > 0$ whenever $a > 0$. The mapping $a \to R(a)$ is monotone and thus continuous on a dense subset of $(0, 1)$. Therefore, without loss of generality we may assume a to be a continuity point of R [if not, just replace a by a continuity point in $(0, a)$]. Using the definition (3.3) of I_2 it holds that

$$\inf_{w \in (\mathcal{N}_a(w^*))^c} I_2(w) = \inf_{w \in ((\mathcal{N}_a(w^*))^c)^\circ} I_2(w) = R(a),$$

and

$$\lim_{T \to \infty} \frac{1}{T} \log P(\rho_T \in (\mathcal{N}_a(w^*))^c) = -R(a).$$

Next, for some $\epsilon > 0$, consider the event

$$\{\eta_T^\infty \in \mathcal{N}_\epsilon(\mu),\ \rho_T \in (\mathcal{N}_a(w^*))^c\}.$$

Using the large deviation upper bound,

$$\begin{aligned} &\limsup_{T \to \infty} \frac{1}{T} \log P\left(\eta_T^\infty \in \mathcal{N}_\epsilon(\mu),\ \rho_T \in (\mathcal{N}_a(w^*))^c\right) \\ &\quad \leq -\inf\left\{I(\gamma, w) : \gamma \in \bar{\mathcal{N}}_\epsilon(\mu),\ w \in (\mathcal{N}_a(w^*))^c\right\}, \end{aligned}$$

where $\bar{\mathcal{N}}_\epsilon(\mu)$ is the closure of $\mathcal{N}_\epsilon(\mu)$. We now claim that for small enough $\epsilon > 0$ the infimum in the last display is strictly larger than $R(a)$. If so, then

$$\begin{aligned} &\limsup_{T \to \infty} \frac{1}{T} \log P\left(\eta_T^\infty \in \mathcal{N}_\epsilon(\mu) | \rho_T \in (\mathcal{N}_a(w^*))^c\right) \\ &\quad \leq -\inf\left\{I(\gamma, w) : \gamma \in \bar{\mathcal{N}}_\epsilon(\mu),\ w \in (\mathcal{N}_a(w^*))^c\right\} + R(a) \\ &\quad < 0, \end{aligned}$$

which gives the exponential decay to zero of the conditional probability.

To show that the infimum is greater than $R(a)$ for $\epsilon > 0$ small enough we argue by contradiction. For the given a the set

$$\begin{aligned} &\{(\gamma, w) : I(\gamma, w) = R(a), w \in (\mathcal{N}_a(w^*))^c\} \\ &\quad = \{(\gamma, w) : I(\gamma, w) \leq R(a), w \in (\mathcal{N}_a(w^*))^c\} \end{aligned}$$

is compact. In addition, by Lemma 5.3 the projection onto the first component does not contain the original invariant measure μ. If the claim is not true then for every $\epsilon = 1/n$, $n \in \mathbb{N}$,

$$\inf\{I(\gamma, w) : \gamma \in \bar{\mathcal{N}}_{1/n}(\mu), w \in (\mathcal{N}_a(w^*))^c\} \leq R(a).$$

Since I has compact level sets this means the infimum is attained, and hence, using compactness of level sets once more, by choosing a convergent subsequence (indexed by n) we have (γ_n, w_n) such that

$$I(\gamma_n, w_n) \leq R(a), \ \gamma_n \to \mu, \text{ and } w_n \to \hat{w} \in (\mathcal{N}_a(w^*))^c.$$

By lower semicontinuity $I(\mu, \hat{w}) \leq R(a) < \infty$. However, this contradicts Lemma 5.3, and completes the proof. ■

6 Further qualitative properties

In this section we study other qualitative properties of the infinite swapping process and, by extension, parallel tempering. In particular, we consider the question of how symmetries and asymmetries of the energy landscape can affect the behavior of infinite swapping. As we will see, when an energy landscape is symmetric in the sense that the energy potential at the minima of the two wells are the same, then subject to the condition that the higher temperature is sufficiently large that the energy barrier is not an obstacle to movement between the wells at that temperature, infinite swapping will converge rapidly. However, when asymmetry holds and the values at the local minima are not the same, then surprising behavior can result. In fact, a "secondary metastability" can emerge, depending on the degree to which the depths of the two wells are not symmetric. This second metastability issue is less of a hindrance than the original energy barrier, but could substantially slow convergence of the weighted empirical measure. In fact a counter-intuitive behavior is observed, in that even *reducing* the energy barrier of one well while holding the other constant may slow the convergence of η_T^∞. Besides a heuristic explanation for this, we will demonstrate the effect via the large deviation rate function using the stochastic control interpretation of the last section.

A second issue we discuss is related to the fact that optimizers in a large deviations analysis can explain how a particular rare event occurs. Suppose one observes that the weighted empirical measure has not properly assigned mass between the wells in such a two well model. By minimizing the rate

function subject to such a constraint (e.g., the well to the right is given a fraction $\kappa(1-\delta)$ of the mass under η_T^∞ when the stationary measure gives it κ), one can find the most likely observed weighted empirical distribution given this "error." The solution to the associated stochastic control problem, and specifically the form of the feedback control, will then identify how this error occurred. In particular, it identifies those places in the state space where poor sampling of the underlying distribution has the largest impact and produces the greatest error. This issue is discussed and illustrated via numerical computation at the end of the section.

6.1 Symmetric and asymmetric double wells

In this subsection and the next the potential will have two local minima at $x_L < 0$ and $x_R > 0$ and a local maximum (top of the separating barrier) at 0. For simplicity we think of the underlying state space $\mathcal{S}$ as being a grid in $\mathbb{R}$ that includes the local minima as well as the top of the barrier. The restriction to $\mathbb{R}$ is to simplify the discussion, while the assumption of a finite state space will allow the explicit numerical solution to certain optimization problems. However, the conclusions will hold more generally within the restriction of a two-well landscape.

We first review certain properties of the INS process. The INS process $\mathbf{Y}$, which takes values in $\mathbb{R}^2$, has four stable points, which one can view as being the local minima of an implied cost potential of the form

$$U(y_1, y_2) = -\log\left(e^{-\frac{1}{\tau_1}V(y_1)-\frac{1}{\tau_2}V(y_2)} + e^{-\frac{1}{\tau_1}V(y_2)-\frac{1}{\tau_2}V(y_1)}\right).$$

Note that regardless of the form of V, this potential is symmetric about the diagonal $y_1 = y_2$, and hence the dynamics are likewise symmetric. The mean value for increments of the INS process are illustrated in Figure 6.1. If one considers $\mathbf{Y}$ as a process in $\mathbb{R}^2$ with the indicated metastable states, then the primary impediment to good sampling and an accurate approximation to the symmetrized distribution $\frac{1}{2}[\mu(x_1, x_2) + \mu(x_2, x_1)]$ is the movement of the process between neighborhoods of the four states $(x_L, x_L), (x_L, x_R), (x_R, x_L)$ and (x_R, x_R).

An alternative perspective is to consider $\mathbf{Y}$ as giving the locations of two particles, whose transition rates are described by (2.4), with ρ from (2.5) of the form

$$\rho(x_1, x_2) = \frac{1}{1 + e^{\left[\frac{1}{\tau_1}-\frac{1}{\tau_2}\right](V(x_1)-V(x_2))}}.$$

Since Γ^1 is the intensity matrix of the low temperature dynamics and Γ^2 that of the high temperature, we see that whenever $V(Y_1)$ is larger than

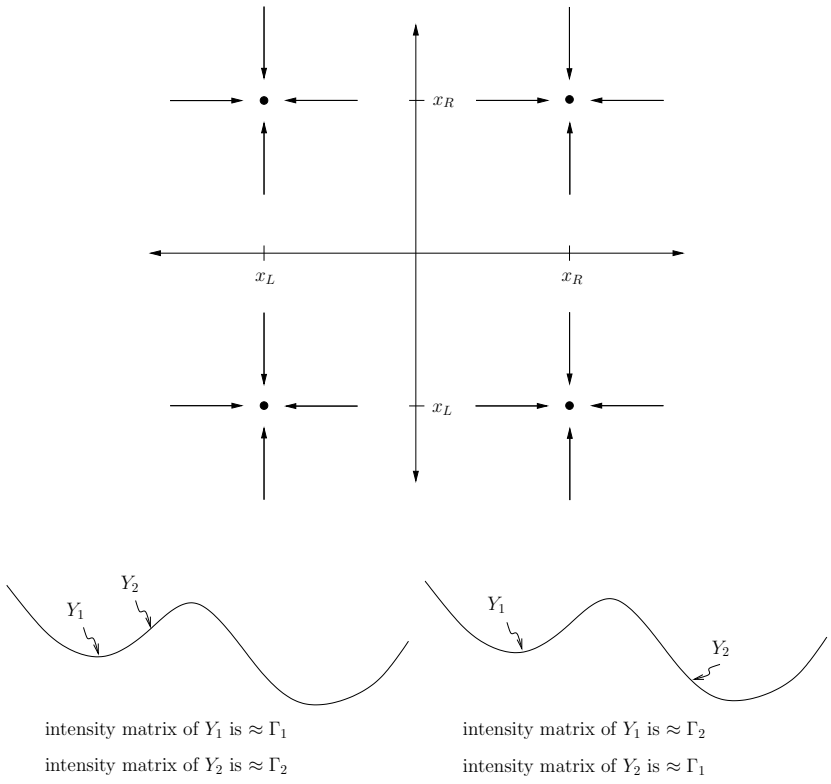

Figure 1: Assignment of dynamics due to relative heights

$V(Y_2)$ by a certain amount, then Y_1 has essentially been given the high temperature dynamics and Y_2 the low, and conversely. This is due to the exponential scaling in ρ, and has an analogue for parallel tempering when the rate of swap attempts is high. The situation is illustrated in Figure 1 for an asymmetric landscape.

We next consider how this "toggling" between high and low temperature dynamics affects the behavior of $\mathbf{Y}$, and in particular how it affects the qualitative properties of the empirical measure of $\mathbf{Y}$ with regard to sampling in $\mathbb{R}^2$. Suppose the well is symmetric as in Figure 2 with wells of depth h. Recall the assumption that the higher temperature is such that the process can easily cross the barrier separating the two wells. We claim that this implies the infinite swapping process $\mathbf{Y}$ easily moves between the four metastable points. Indeed, if both particles are placed in the left well [so that $\mathbf{Y}$ is near (x_L, x_L)] then after a relatively short time one of the two

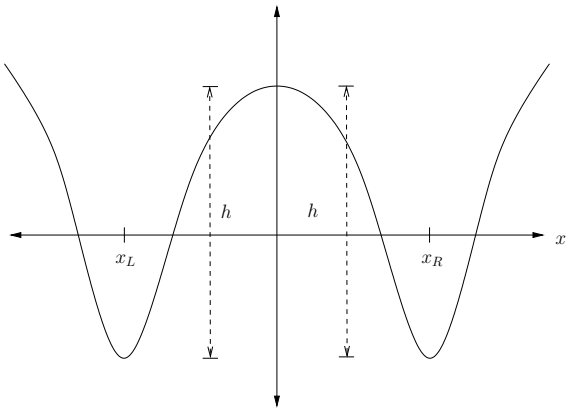

Figure 2: Symmetric double well

particles will scale the landscape to a point somewhere between x_L and 0, while the other is near x_L. At this stage the particle higher on the energy landscape essentially follows the dynamic of the intensity matrix Γ_2, at least until it and the particle near x_L reach places in the landscape of the same height. This can happen by the particle falling back to the basin of the left well, or by crossing the barrier to enter the neighborhood of x_R. By assumption, the latter is not a particularly rare event, and corresponds to $\mathbf{Y}$ ending up near (x_L, x_R) or (x_R, x_L). The argument can be repeated and, owing to the symmetry, movement of $\mathbf{Y}$ between neighborhoods of all the points $(x_L, x_L), (x_L, x_R), (x_R, x_L)$ and (x_R, x_R) occurs with the same frequency.

We next contrast this with what can be expected when the heights are changed, and in particular if there is a reduction of the height of the right well, so that in Figure 3 we have $h_L = h$ and $h_R < h$. In the case the discussion just given applies equally well when both particles start in a neighborhood of x_L, but the behavior of the process now changes radically when $\mathbf{Y}$ is near (x_L, x_R) or (x_R, x_L). To be specific, assume that it is near (x_L, x_R). The transitions of interest are: (a) Y_1 joins Y_2 in the right well, and (b) Y_2 returns and joins Y_1 in the left well. The event (b) is actually quite likely, since if Y_1 is in the deeper well then it is highly probable that it is lower on the energy landscape, and therefore Y_2 is given the high temperature dynamic. For the same reason (a) is unlikely. Indeed, the only way it can happen is if Y_1, in spite of being given the lower temperature dynamics, is able to move up the landscape to a point were it exceeds the typical energy value that Y_2 sees while in the right well (and using the higher temperature

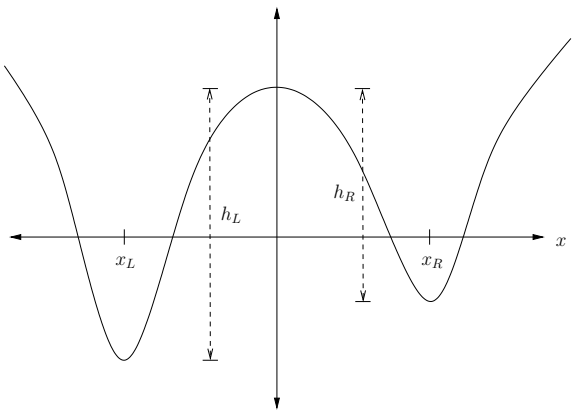

Figure 3: Asymmetric double well

dynamic). Thus Y_1 must overcome an energy barrier, whose size is related to the degree of asymmetry of the two well depths. We call this a *secondary metastability,* and note that the effect of lowering one of the energy barriers in the single particle model is that it increases an energy barrier for the two particle INS process, leading to poorer sampling of the state space by the process.

6.2 Numerical example

The issue described in the last section is reflected in the large deviation rate function. This will be illustrated by numerically solving a constrained optimization problem, though with a constraint of a different form from the last section. We are particularly interested in the impact of the secondary metastability on the accuracy of integrals with respect to the low temperature marginal. Specifically, we consider

$$\inf_{\nu\in\mathcal{P}(\mathcal{S}^2)}\left\{J(\nu):\sum_{\mathbf{x}\in[0,\infty)\times\mathcal{S}}\rho(\mathbf{x})\nu(\mathbf{x})+\rho(\mathbf{x}^R)\nu(\mathbf{x}^R)=\kappa(1-\delta)\right\},$$

where as before $\mathbf{x}^R = (x_2, x_1)$. Here $\kappa \in (0,1)$ is the mass that the low temperature marginal places on the set $[0,\infty)$, and δ is the size of the error. Recalling that $\rho(\mathbf{x})\nu(\mathbf{x}) + \rho(\mathbf{x}^R)\nu(\mathbf{x}^R)$ is the mapping that takes a symmetrized measure to its weighted counterpart, this variational problem will identify the most likely distribution for the mass given that it is incorrectly

assigned by the sampling. Note that one expects the relative distribution within each well to converge much faster than the relative weights between wells. Hence this constraint focuses attention on the most likely error that the sampling must overcome, which is properly assigning the mass between the two wells.

We study this problem using the same methods as in the previous section, i.e., introduce a Lagrange multiplier and analyze the resulting ergodic control problem. Since the only difference is the form of the constraint, we do not repeat any details in the derivation of the control problem. For numerical illustration of the effect of asymmetry in the potential landscape we use the form of Glauber dynamics for two temperatures as defined in Section 2. The claim of the previous subsection, based on the heuristic discussion involving the stochastic control problem of Section 4, is that the secondary metastability induced by an asymmetric potential V slows the convergence of η_T^∞ to μ. To demonstrate this effect, we show that the rate associated with the same relative error (i.e., value of δ) is lower for the asymmetric case, indicating that the outcome is more likely.

In order to have a convenient way of constructing potentials with similar shape but varying degree of asymmetry we use the following family of functions, sometimes referred to as Franz potential,

$$V(x) = V(x;\alpha) = \frac{3x^4 - 4(\alpha-1)x^3 - 6\alpha x^2}{2\alpha+1} + 1,$$

For every value of α, $V(\cdot;\alpha)$ has a fixed local minimum at $x_L = -1$, a varying local minimum at $x_R = \alpha$ and a fixed barrier of height 1 at the origin. Figure 4 shows the potential V for some values of α. In particular, taking $\alpha = 1$ produces a symmetric two well potential and $\alpha = 0$ produces a single well.

Table 1 shows the value of $\kappa = \mu_1([0,\infty))$ for some values of α when the Franz potential is used to define the underlying Gibbs measure.

α	1	0.97	0.95	0.90	0.85
κ	0.500	0.318	0.223	0.0840	0.0316

Table 1: The probability $\kappa = \mu_1([0,\infty))$ for some values of α in the Franz potential

In Table 2 numerical results for the discussed optimization are presented. Table 3 repeats the results but now normalized to the symmetric case $\alpha = 1$ for each δ. The results in Tables 2-3 illustrate how the rate function changes with κ and δ when the measures low-temperature marginal is restricted to

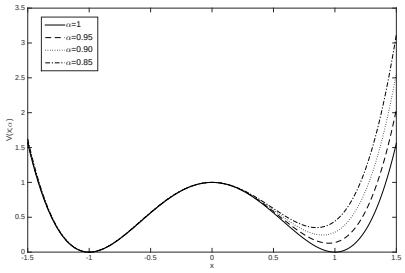

Figure 4: Franz potential V for some values of α.

	α				
δ	1	0.97	0.95	0.90	0.85
0.05	1.5250e-5	8.5478e-6	6.0461e-6	2.7959e-6	1.4012e-6
0.10	6.1151e-5	3.4911e-5	2.4887e-5	1.1609e-5	5.8911e-6
0.15	1.3802e-4	8.0513e-5	5.7975e-5	2.7562e-5	1.4163e-5
0.20	2.4655e-4	1.4704e-4	1.0702e-4	5.1900e-5	2.7206e-5

Table 2: Large deviation rate for the minimizing measure with a low-temperature marginal that puts mass $\kappa(1-\delta)$ in the shallow well, for different values of α in the Franz potential V; $\tau_1 = 0.1,\ \tau_2 = 0.5,\ |\mathcal{S}| = 12$.

assign less mass to $[0, \infty)$ than the true invariant distribution $\bar{\mu}$. In particular, it illustrates how the optimal value of the rate function decreases with the level of asymmetry when the amount of mass that is redistributed is fixed. This observation corresponds precisely to the large deviation interpretation that for increased level of asymmetry the empirical measure will take a longer time to converge.

6.3 How poor sampling will occur

Our second and final use of the rate function to study qualitative properties of Monte Carlo addresses the following question. Suppose that a given Markov process has invariant distribution μ that is concentrated in two wells (as in a single temperature version of the model just considered). Suppose we also consider a measure other than μ that one is likely to see prior to convergence, e.g., a minimizer of the rate function subject to a constraint

	α				
δ	1	0.97	0.95	0.90	0.85
0.05	1	0.5605	0.3965	0.1833	0.09188
0.10	1	0.5709	0.4070	0.1898	0.09634
0.15	1	0.5833	0.4200	0.1997	0.1026
0.20	1	0.5964	0.4341	0.2105	0.1103

Table 3: Large deviation rate for the minimizing measure with a low-temperature marginal that puts mass $\kappa(1-\delta)$ in the shallow well, normalized to the rate for $\alpha = 1$ (symmetric potential); $\tau_1 = 0.1,\ \tau_2 = 0.5,\ |\mathcal{S}| = 12$.

on improperly assigning mass to the two wells (also as in the last section). By solving the associated stochastic control problem we will find the change of measure (change of jump rates) which minimizes the average cost per unit time to hit the given target measure. Using the large deviation rate as in the proof of Proposition 5.1 to bound conditional probabilities given a certain outcome, one can characterize the parts of the state space where the observed empirical data collected along the simulated trajectory deviates from what is expected based on the underlying dynamics. In other words, if one were to attempt to infer the true dynamics based on the empirical data, the solution to the control problem will tell us where these inferred dynamics will deviate most from the true dynamics. One can imagine that is at precisely these locations that greater accuracy in sampling (e.g., in approximating transition probabilities) would have the greatest impact on overall performance. Although we do not propose a particular use along those lines at this time, it seems to be interesting information with some potential for improving schemes.

Here we consider the case of only one temperature $\tau = 0.1$ and the optimization problem

$$\inf_{\nu \in \mathcal{P}(\mathcal{S})} \left\{ J(\nu) : \nu([0,\infty)) = \kappa(1-\delta) \right\},$$

where κ is the amount of mass the invariant distribution of the underlying Glauber dynamics - with Franz potential - puts in the right well. As in the previous subsection, the constraint amounts to placing less mass in the shallow well compared to the invariant measure μ_1.

Figure 5 shows the value function W that is associated with the solution ν to the optimization problem for $\alpha = 1$ (symmetric well). From the value function we can compute the additional factor $\exp\{W(y) - W(x)\}$ in the optimal control for a jump from x to y; Figures 6 and 7 show the extra

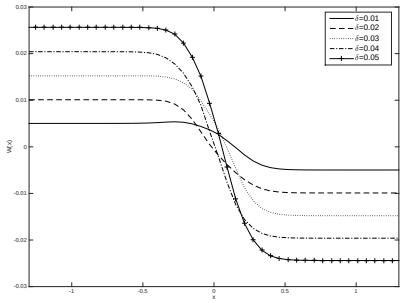

Figure 5: Value function W when $\alpha = 1$; $|\mathcal{S}| = 50$, $\tau = 0.1$.

factor when y is one step to the right and left, respectively. Note that the controlled jump rates differ from the uncontrolled ones in a neighborhood of the origin. Of course these states are not visited much during the simulation of a trajectory, and for this reason alone one might expect the numerics to poorly approximate the true distributions (e.g., point to point conditional densities). However, this statement applies to many parts of the state space, and the critical difference is that errors here are more important in producing error to the approximation of the invariant distribution. A possible fix would be to spend some computational effort estimating the critical quantities (e.g., the conditional probability to send at ± 0.25 after reaching 0.0) beforehand, and then use these more accurate estimates as the basis for a simulation schemes that excises the corresponding parts of the simulated trajectory.

A Ancillary results

Lemma A.1 *For any sequence $a_1, a_2, \ldots, a_K$ such that $a_i \geq 0$ for all i and $K \in [0, \infty)$,*

$$\sum_{i=1}^{K} a_i = K$$

and

$$\sum_{i=1}^{K} a_i^{1/2} > K$$

cannot both be true.

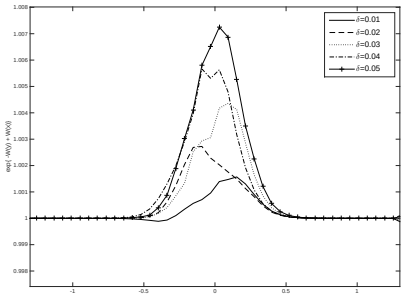

Figure 6: Additional factor in optimal control when $\alpha = 1$ and jumps to the right; $|\mathcal{S}| = 50$, $\tau = 0.1$.

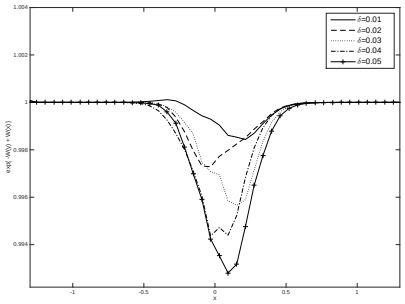

Figure 7: Additional factor in optimal control when $\alpha = 1$ and jumps to the left; $|\mathcal{S}| = 50$, $\tau = 0.1$.

Proof. We can assume without loss that $K > 0$. Let $b_i = a_i/K$, so that $\{b_i, i = 1, \dots, K\}$ is a probability. By Hölder's inequality

$$\sum_{i=1}^{K} b_i^{1/2} \le K^{1/2} \left(\sum_{i=1}^{K} b_i \right)^{1/2} = K^{1/2}.$$

Using $b_i = a_i/K$ gives $\sum_{i=1}^{K} a_i^{1/2} \le K$, which completes the argument. ■

Lemma A.2 *Consider the ergodic control problem or equivalent minimization problem of Lemma 5.4, with $h(\mathbf{x}) = \lambda\rho(\mathbf{x})$. For two temperatures the optimal cost γ^* satisfies*

$$\gamma^* < \frac{\lambda}{2}.$$

In the general case with K temperatures, $\gamma^* < \lambda/K$!

Proof. To simplify the notation we consider the case $\lambda = 1$. From Lemma 5.4 we know that there is a minimizing measure ν^* in (5.1) and that the optimal cost γ^* satisfies

$$\gamma^* = J(\nu^*) + \sum_{\mathbf{x} \in \mathcal{S}^2} \rho(\mathbf{x}) \nu^*(\mathbf{x}).$$

Moreover, if W^* is defined as

$$W^*(\mathbf{x}) = -\log \left[\frac{d\nu^*}{d\bar{\mu}} \right]^{1/2} (\mathbf{x}), \quad \mathbf{x} \in \mathcal{S}^2,$$

then (γ^*, W^*) is a solution to the Bellman equation (4.9).

Suppose that W^* is a constant. Inserting this into the Bellman equation yields, for each $\mathbf{x} \in \mathcal{S}^2$,

$$0 = -\gamma^* + \rho(\mathbf{x}),$$

which cannot hold since ρ is not a constant. Thus, W^* cannot be a constant function. This in turn implies that the likelihood ratio $[d\nu^*/d\bar{\mu}]$ is not constant equal to 1 (the only possible constant value). Thus ν^* is not $\bar{\mu}$, the invariant measure for the original symmetrized dynamics.

Inserting $\bar{\mu}$ into the objective function in (5.1) gives

$$J(\bar{\mu}) + \sum_{\mathbf{x} \in \mathcal{S}^2} \rho(\mathbf{x}) \bar{\mu}(\mathbf{x}) = \sum_{\mathbf{x} \in \mathcal{S}^2} \rho(\mathbf{x}) \bar{\mu}(\mathbf{x}) = \frac{1}{2},$$

where the second equality comes from $\rho(\mathbf{x})+\rho(\mathbf{x}^R)=1=1/2+1/2$ and the symmetry of $\bar{\mu}$. Thus the cost associated with the uncontrolled dynamics is $1/2$. Since ν^* is the unique minimizer in (5.1) and $\nu^*\neq\bar{\mu}$, it holds that

$$\gamma^*=J(\nu^*)+\sum_{\mathbf{x}\in\mathcal{S}^2}\rho(\mathbf{x})\nu^*(\mathbf{x})<J(\bar{\mu})+\sum_{\mathbf{x}\in\mathcal{S}^2}\rho(\mathbf{x})\bar{\mu}(\mathbf{x})=\frac{1}{2}.$$

The argument for $K>2$ temperatures is completely analogous. ■

Lemma A.3 *Assume S is a finite set and that $\Gamma_{x,y}$, $x,y\in S$ is the intensity matrix of an ergodic Markov chain on S with invariant probability distribution $\bar{\mu}$. Let $q(x)=\sum_{y\in S}\Gamma_{x,y}$, and for $\nu\in\mathcal{P}(S)$ with $\theta(x)=\nu(x)/\bar{\mu}(x)$ let*

$$J(\nu)=\sum_{x\in S}q(x)\theta(x)\bar{\mu}(x)-\sum_{x,y\in S}\theta^{1/2}(x)\theta^{1/2}(y)\Gamma_{x,y}\bar{\mu}(x).$$

Then $J(\nu)$ is strictly convex on the relative interior of $\mathcal{P}(S)$.

Proof. It is enough to show the strict convexity of

$$\theta(\cdot)\to-\sum_{x,y\in S}\theta^{1/2}(x)\theta^{1/2}(y)\Gamma_{x,y}\bar{\mu}(x)$$

for $\theta(x)\geq 0$, $\sum_{x\in S}\theta(x)\bar{\mu}(x)=1$. Let $\{x_1,x_2,\ldots,x_K\}$ be an enumeration of the distinct elements of S, $\theta_i=\theta(x_i),\bar{\mu}_i=\bar{\mu}(x_i)$ and $f_{i,j}(\theta)=-\theta_i^{1/2}\theta_j^{1/2}$. If $M_{i,j}(\theta)$ denotes the matrix of second order partial derivatives of $f_{i,j}(\theta)$ at θ, then straightforward calculation shows that the eigenvalue 0 is repeated $K-1$ times, and $(\theta_i/\theta_j+\theta_j/\theta_i)$ is also an eigenvalue with eigenvector $\theta_j e_i-\theta_i e_j$. Hence the null space of this matrix is the collection of vectors orthogonal to $\theta_j e_i-\theta_i e_j$. Since $(\theta_i/\theta_j+\theta_j/\theta_i)>0$, $f_{i,j}(\theta)$ is strictly convex (as a function in $\mathbb{R}^K$) at θ except in those directions orthogonal to $\theta_j e_i-\theta_i e_j$.

Since $\Gamma_{x,y}$ is ergodic all states communicate, and so there exists a sequence $1=i_1,i_2,\ldots,i_K,i_{K+1}=1$ such that $\Gamma_{x_{i_k},x_{i_{k+1}}}>0$ for $k=1,\ldots,K$. Thus $-\sum_{x,y\in S}\theta^{1/2}(x)\theta^{1/2}(y)\Gamma_{x,y}\bar{\mu}(x)$ is strictly convex except in those directions that are orthogonal to each of $\theta_{i_{k+1}}e_{i_k}-\theta_{i_k}e_{i_{k+1}}$, which is exactly the set of directions spanned by $(\theta_1,\theta_2,\ldots,\theta_K)$. Since this direction cannot be parallel to $\{\theta:\sum_{k=1}^K\theta_k\bar{\mu}_k=1\}$, $J(\nu)$ is strictly convex on this set. ■

Remark A.4 The proofs in Section 5 were largely confined to the setting of two temperatures τ_1,τ_2. This was to keep the notation simple and the results generalize to any number $K\geq 2$ of temperatures. The only result

which appears to substantially use that two temperatures are considered is Lemma 5.3, and specifically the argument by contradiction. Here we outline how the proof would proceed for the general setting.

In the setting of K temperatures the assumption (5.7) becomes

$$\sum_{\sigma\in\Sigma} e^{-2W(\mathbf{x}^\sigma)} - K! = 0, \ \forall \mathbf{x} \in \mathcal{S}^K. \tag{A.1}$$

We still have that $W(\mathbf{x}) = 0$ for $\mathbf{x} \in \mathcal{D}$.

Let $\mathcal{D}$ denote the set of diagonal states: $\mathcal{D} = \{\mathbf{x} : \mathbf{x} = \mathbf{x}^\sigma, \ \forall \sigma \in \Sigma_K\}$. The only such states are those for which all components are equal. The cost structure is such that $h(\mathbf{x}) = 1/K!$ for $\mathbf{x} \in \mathcal{D}$.

Consider the states that communicate directly with $\mathcal{D}$, i.e., those only one step away from a diagonal state. Since the underlying processes only jump one at a time there can only be a difference in one component, the others remaining fixed. There are a total of $K!$ possible permutations in Σ_K, $(K-1)!$ of which keep a specific component fixed. Thus, for a state that is one step removed from the diagonal there are $(K-1)!$ permutations that result in the same state. Moreover, the diagonal state in question will communicate directly with the remaining K permutations as well.

The states one step away from a specific diagonal point can be viewed as forming disjoint sets of states according to the previous description. For each state $\mathbf{y}$ one step removed from an $\mathbf{x}$, there are K distinct states $\mathbf{y}_1, \ldots, \mathbf{y}_K$ that are permutations of $\mathbf{y}$ and communicate directly with $\mathbf{x}$. For each such collection of states we can pick one to represent the collection (does not matter which one we pick). Let $\mathcal{A}_x$ denote the collection of such representative states $\mathbf{y}$. In the case of two temperatures this can be phrased as only looking at states above the diagonal.

The Bellman equation for a diagonal state $\mathbf{x}$ takes the form

$$0 = \sum_{\mathbf{y}\in\mathcal{A}_x} \sum_{\sigma:\mathbf{y}^\sigma\neq\mathbf{y}} r(\mathbf{x},\mathbf{y}) \left[1 - e^{-W(\mathbf{y}^\sigma)+W(\mathbf{x})}\right] - \gamma + \frac{1}{K!}.$$

The rates $r(\mathbf{x}, \mathbf{y}^\sigma)$ are all equal due to symmetry. Combined with $W(\mathbf{x}) = 0$ for $\mathbf{x} \in \mathcal{D}$, this allows the Bellman equation to be expressed as

$$0 = \sum_{\mathbf{y}\in\mathcal{A}_x} r(\mathbf{x},\mathbf{y}) \left[K - \sum_{\sigma:\mathbf{y}^\sigma\neq\mathbf{y}} e^{-W(\mathbf{y}^\sigma)}\right] - \gamma + \frac{1}{K!}.$$

Since $\gamma < (1/K!)$ and the rates are all non-negative it must be the case that

for at least one $\mathbf{y} \in \mathcal{A}_x$

$$\sum_{\sigma:\mathbf{y}^\sigma \neq \mathbf{y}} e^{-W(\mathbf{y}^\sigma)} - K > 0.$$

For states one step from the diagonal, since $(K-1)!$ permutations will result in the same state, the condition (A.1) takes the form

$$(K-1)! \sum_{\sigma:\mathbf{x}^\sigma \neq \mathbf{x}} e^{-W(\mathbf{x}^\sigma)} - K! = 0 \Leftrightarrow \sum_{\sigma:\mathbf{x}^\sigma \neq \mathbf{x}} e^{-W(\mathbf{x}^\sigma)} - K = 0.$$

That is we need only be concerned with the permutations that switch the location of the component that differs from the diagonal state (and the σ that corresponds to the identify map in Σ_K). There will then be $(K-1)!$ permutations that produces the exact same state, yielding the factor $(K-1)!$ in front of the sum.

For the reduced form of (A.1) and the Bellman equation to hold, we must have that

$$\sum_{\sigma:\mathbf{x}^\sigma \neq \mathbf{x}} e^{-W(\mathbf{x}^\sigma)} - K = 0,$$

for all $\mathbf{y}$ that communicate with $\mathbf{x}$, and for at least one such $\mathbf{y}$,

$$\sum_{\sigma:\mathbf{y}^\sigma \neq \mathbf{y}} e^{-W(\mathbf{y}^\sigma)} - K > 0.$$

This is precisely the setting of Lemma A.1 with the a_is represented by $e^{-2W(\mathbf{y}^\sigma)}$ for the K relevant permutations σ. The lemma then implies that (A.1) is inconsistent with the Bellman equation and therefore cannot hold. This contradicts that $(M\bar{\nu}) = \mu$.

References

[1] J. Doll and P. Dupuis. On performance measures for infinite swapping Monte Carlo methods. *J. Chem. Phys*, 142:024111, 2015.

[2] J. Doll, N. Plattner, D. L. Freeman, Y. Liu, and P. Dupuis. Rare-event sampling: occupation-based performance measures for parallel tempering and infinite swapping Monte Carlo methods. *J. Chem. Phys*, 137:204112, 2012.

[3] P. Dupuis and R. S. Ellis. Large deviation analysis of queueing systems. In F. P. Kelley and R. J. Williams, editors, *Stochastic Networks*, volume 71 of *IMA Volumes in Mathematics and Its Applications*, pages 347–356. Springer–Verlag, New York, 1995.

[4] P. Dupuis and R. S. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*. John Wiley & Sons, New York, 1997.

[5] P. Dupuis and Y. Liu. On the large deviation rate for the empirical measure of a reversible pure jump markov processes. *Annals of Probability*, 43:1121–1156, 2015.

[6] P. Dupuis, Y. Liu, N. Plattner, and J.D. Doll. On the infinite swapping limit for parallel tempering. *SIAM J. Multiscale Model. Simul.*, 10:986–1022, 2012.

[7] D. J. Earl and M. W. Deem. Parallel tempering: theory, applications, and new perspectives. *Phys. Chem. Chem. Phys.*, 7:3910–3916, 2005.

[8] S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, New York, 1986.

[9] W. H. Fleming. Exit probabilities and optimal stochastic control. *Applied Math. and Optimization*, 4:329–346, 1978.

[10] W. H. Fleming and H. M. Soner. *Controlled Markov processes and viscosity solutions* Springer-Verlag, New York, 1993.

[11] C. J. Geyer. Markov chain Monte Carlo maximum likelihood. In *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, pages 156–163, New York, 1991. American Statistical Association.

[12] H. J. Kushner and P. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, New York, 2001. Revised Second Edition.

[13] D. G. Luenberger. *Optimization by Vector Space Methods*. John Wiles & Sons, Inc., New York, NY, first edition, 1969.

[14] L. C. G. Rogers and D. Williams. *Diffusions, Markov processes, and martingales. Vol. 2*. Cambridge Mathematical Library. Cambridge University Press, Cambridge, 2000. Itô calculus, Reprint of the second (1994) edition.

[15] D. W. Stroock. *An Introduction to Markov Processes*, volume 230 of *Graduate Texts in Mathematics*. Springer, Berlin, 2005.

[16] Y. Sugita and Y. Okamoto. The incomplete beta function law for parallel tempering sampling of classical canonical systems. *Chem. Phys. Lett.*, 314:141–151, 1999.

[17] R. H. Swendsen and J. S. Wang. Replica Monte Carlo simulation of spin glasses. *Phys. Rev. Lett.*, 57:2607–2609, 1986.